VII

An Actuarial Layman's Guide to Building Stochastic Interest Rate Generators

James A. Tilley

Abstract

A stochastic interest rate generator is a valuable actuarial tool. The parameters that specify a stochastic model of interest rates can be adjusted to make the model arbitrage-free, or they can be adjusted to accommodate an individual investor's subjective views. The arbitrage-free settings of the parameters must be used when pricing streams of interest-rate-contingent cash flows, for example, when establishing the risk-neutral position for asset-liability management. The real-world settings of the parameters should be used when evaluating the risk-reward tradeoffs inherent in deviating from the risk-neutral position.

Without relying on formulas, this paper presents the important concepts underlying the theory of arbitrage-free pricing of interest-rate-contingent cash flows: absence of opportunities for riskless arbitrage; completeness of markets; relative prices that do not depend on individual investors' subjective views or risk preferences; and expected-value pricing in the risk-neutral world. Using these concepts, the paper then describes the steps required to build continuous stochastic models of interest rates, including models that are either partially or fully arbitrage free. After studying the paper, all actuaries should be able to comprehend better some of the literature in this important subject area. Then, after studying some of the technical references, many actuaries should be in a position to begin to build their own practical models.

1. Introduction

In recent years, the literature of financial economics has featured papers on how to value interest-rate-contingent claims by means of option-pricing models (for example, [2], [8], [11], [14], and [18]). The most important applications include the pricing of fixed-income instruments with embedded options: callable corporate bonds, mortgage-backed securities subject to prepayment risk, collateralized mortgage obligations (CMOs) created by allocating the cash flows arising from pools of mortgages to different classes of bonds, floating-rate and other indexed bonds, and various hedging instruments, such as futures, options, and interest rate swaps, caps, and floors. Life insurers have begun to use option-pricing models to value complicated interest-rate-contingent liabilities that contain embedded options, such as the insurer's right to reset periodically the interest rate credited to a policyholder's account, or the policyholder's right to take loans at below-market interest rates or to surrender a policy for a cash amount that does not fully take into account the level of interest rates prevailing at the time of surrender. In the U.S., the capital adequacy of depository institutions (banks and thrifts) is now measured against risk-based capital guidelines that include an interest rate risk component, for which an option-pricing model is needed to value mortgage-related assets properly.

Unfortunately, the papers about option pricing are often very technical, leaving almost all actuaries frustrated, because they recognize the importance of utilizing option-based models, but they do not understand the theory well enough to be able to write computer programs to implement it. I have made no attempt in this paper to review the literature on the subjects of interest rate models and option pricing. That would have diluted my efforts in achieving the paper's objective of bringing the actuary who is not an expert in either financial economics or in the mathematics of stochastic processes (martingales and stochastic calculus, in particular) comfortably to the point of understanding how a useful model for valuing streams of interest-rate-contingent cash flows can be built. Several routes could have been followed to achieve this objective. After much consideration, I decided to develop the paper without formulas, with one exception: in offering an example of a continuous stochastic process for interest rates, it seemed easier to write down a few equations than to write elaborately around them. After reading this paper, and perhaps relying to some extent on the references cited, mathematically inclined actuaries will likely be able to construct stochastic interest rate generators appropriate to their needs. Other actuaries, if unable to build such generators themselves, should at least be able to apply the generators in solving asset and liability valuation problems. The principal goal of this paper is to discuss thoroughly the concepts underlying the valuation of interest-rate-contingent streams of cash flows, not to provide a set of mathematical recipes that can be programmed into an option-pricing model on a computer.

Central to the problem of valuing interest-rate-contingent cash flows is the creation of an appropriate set of interest rate paths or scenarios. In fact, once a theoretically sound stochastic interest rate generator has been constructed, all the applications described above can be handled. Each application involves projecting the relevant cash flows along a path, then discounting the projected cash flows for the path, using the short-term interest rates along the path, to a present value number for the path, and finally averaging the present value numbers for all paths to obtain the arbitrage-free value of the cash-flow stream. The rigorous proof that such a simple procedure works is highly mathematical (see, for example, the texts [12] and [16]). However, one can develop an intuitive feel for the validity of the approach without having to face intimidating mathematics. In this paper, I offer some explanation that serves to build such intuition, but not so much as to dis-

tract us from the main goal of laying the foundation for constructing arbitrage-free stochastic interest rate generators.

Section 2 introduces the concepts of current-coupon yields, spot rates of interest, and forward rates of interest, and describes the relationships among them. Section 3 briefly describes both discrete-state and continuous-state models of interest rates and debates the strengths and weaknesses of each. Section 4 introduces several key concepts from financial economics, and then indicates how the assumptions of complete markets and the lack of riskless arbitrage opportunities allow one to move into a special equilibrium world characterized as *risk neutral,* in which the valuation of interest-rate-contingent cash flows becomes a straightforward expected-value problem. Section 5 fulfills the purpose of the paper by documenting how to construct a path generator based on a continuous process, and Section 6 then indicates how such a generator can be used. Section 7 lists the key conclusions of the paper.

2. Yield Curve and Term Structure

This paper focuses on interest rates for instruments free from default and call risk, which, in the financial markets in the U.S., means U.S. Treasury bills, notes, and bonds. All other investment-grade fixed-income financial assets are priced relative to U.S. Treasury obligations. There are several equivalent ways to express the set of yields applying to risk-free debt obligations of various maturities. The most common, because it is the basis on which traders make quotations, is the concept of the *yield curve.* The yield curve is a graph that depicts the yields of hypothetical U.S. Treasury obligations that trade at a price of par as a function of their terms to maturity in years. By convention, the yields on such *par* bonds are expressed as annual rates of interest payable semiannually, referred to as *bond-equivalent yields,* because bonds issued in the U.S. usually pay coupons semiannually. The hypothetical bonds trading at a price of par that constitute the yield curve are said to have *current coupons.*

Another way to express the information contained in the yield curve is to compute the yields of *zero-coupon* bonds of various maturities from the yields of all current-coupon bonds. A unit par value zero-coupon bond having a maturity of n years pays its holder $1 at the end of n years and nothing before then. A zero-coupon bond is sometimes referred to as a *pure discount* bond, because it must always trade at a price less than par, that

is, at a discount to par. The yield of a zero-coupon bond with maturity n years is referred to as the n-year *spot yield* or *spot rate*. The graph that depicts default-free spot rates as a function of term to maturity is known as the *term structure* of interest rates. The prices of the zero-coupon bonds that define the term structure are often referred to as *spot prices* and, as already stated, are always less than par.

Yet a third way to express the information contained in either the yield curve or the term structure is to compute the yields for *forward* loans. For example, an investor might agree to lend a borrower money in m years and to be repaid in full (principal plus all accumulated interest) in n years from that point in time—that is, at the end of $m + n$ years from today. Such an arrangement is known as an m-year forward n-year loan. The rate of interest for such a loan is referred to as the (m,n) year *forward rate*. More generally, the $(m,n;t)$ year forward rate refers to the interest rate on a loan that will be arranged t years from today, under which an investor will lend a borrower money $m + t$ years from today and will be repaid in full $m + n + t$ years from today.

Using the terms defined above, it can be shown that the n-year spot price is equal to the product of n positive discount factors. The first factor involves only the $(0,1)$ year forward rate; the second factor involves only the $(1,1)$ year forward rate; and the n-th factor involves only the $(n - 1,1)$ year forward rate. Thus, it follows that the $(n + 1)$-year spot price divided by the n-year spot price is equal to a positive number that depends only on the $(n,1)$ year forward rate. This number will be less than or equal to 1 (in other words, it will be a "discount" factor) if, and only if, the $(n,1)$ year forward rate is non-negative.

The information contained in the sets of current-coupon yields, spot rates, and forward rates is equivalent. (Further material on this subject can be found in the text by Sharpe and Alexander [22].) Any one set of yields or rates is sufficient to derive the other two sets. Depending on the situation, there may be a natural set to use, but all carry identical information. For example, when speaking with traders or portfolio managers to obtain interest rate assumptions for pricing an annuity product, an actuary would likely ask about the yield curve. When discounting a stream of fixed and certain cash flows arising from structured settlement annuity liabilities to obtain a current market value, an actuary would naturally use spot rates. When constructing an arbitrage-free theory of interest rate dynamics, most financial economists would use forward rates as the starting point.

3. Discrete versus Continuous Models

Throughout this paper, the term *state of the world* refers to the yield curve prevailing at a particular time or epoch. In a model of interest rates, the adjectives *discrete* and *continuous*, without any modifiers, are used best to describe the type of states of the world represented, not the type of time interval used. In practical applications, regardless of the model used, cash flows are assumed to occur at discrete time intervals: monthly for typical mortgages; quarterly for CMOs, preferred stocks, and some floating-rate bonds; and semiannually for typical bonds. In asset-liability cash-flow analyses, quarter-year periods typically are used. So the basic issue is not whether discrete-time models are to be preferred to continuous-time models; rather, it is whether discrete-state models are to be preferred to continuous-state models.

Most of the recent literature describes discrete models, in which the states of the world are represented by nodes on a lattice (refer to the papers [2], [11], and [18] cited earlier). The vast majority of such models utilize binomial lattices, on which the world evolves from any given state at one epoch to one of two different states at the next epoch. These two states at the end of a time interval are usually referred to as the *up* state and the *down* state with respect to the state at the beginning of the interval. For reasons of computational efficiency, *connected* lattices are almost always used. From any node in a connected lattice, the two-period evolution of states *up first, then down* and the two-period evolution of states *down first, then up* must lead to the same ending node. In a connected lattice, the world can evolve from a single initial state at epoch 0 to one of two states at epoch 1, to one of three states at epoch 2, and so on, to one of $H + 1$ states at epoch H. In a connected binomial lattice model, it is unlikely that the possible states of the *real world* will be sampled sufficiently finely at the early epochs. To remedy this problem, the time interval can be reduced. For example, with daily intervals, there are about 30 states at the end of any one-month period, but the computational demands of creating and using such a model can be enormous, especially for long-term assets or liabilities. Moreover, it is unnatural (and should be unnecessary) to choose a time interval much shorter than the shortest period between cash flows for typical assets and liabilities. Thus, the coarseness-of-sampling difficulty of connected lattice models

remains in many practical situations. Continuous models do not suffer this weakness.

Continuous-state models are described in the academic literature by means of differential equations that represent continuous-time stochastic processes (refer to the papers [8] and [14] cited earlier, and also to the text by Hull [13]). For practical applications, though, the continuous-time process needs to be sampled only at regular time intervals, and the models are reformulated as stochastic difference equations. The time interval is often chosen to equal the shortest period between the cash flows for the assets and liabilities under study. In continuous models, one samples *paths* of interest rates by iterating the difference equation. If P interest rate paths are used, there are P states of the world represented at every epoch. Because the sample of P states at each epoch is drawn from a continuous distribution, the resulting paths of interest rates do not appear to have been constructed artificially. Stated a little differently, it is difficult for an experienced portfolio manager to tell whether an interest rate path was generated from a good continuous model or was constructed from segments of actual interest rate history. The same claim cannot be made for interest rate paths sampled from a lattice.

A connected lattice model has a significant weakness that can be overcome by using a continuous model. For a connected lattice to be *arbitrage-free* (defined in Section 4), severe constraints have to be placed on how it is constructed. These constraints greatly limit the possible yield curve dynamics, and for most models, the resulting evolution of yield curves does not correspond adequately to real-world behavior. The problem arises in simple lattice models because a *single* stochastic factor—the short-term rate of interest—drives the dynamics of the *entire* yield curve, resulting in perfect correlation of yield movements across the curve. In the real world, the movements of neighboring segments of the yield curve may be highly correlated, but they are not perfectly correlated. Arbitrary correlation can be accommodated in a continuous model, because different parts of the yield curve can be assumed to follow correlated stochastic processes.

Discrete and continuous models can also be compared for computational efficiency, which depends on the type of problem to be solved. In the case of interest-rate-contingent, but *path-independent*, cash flows, as are usually associated with pure options, callable bonds, and optional sinking fund bonds, *backward induction* algorithms can be used on a lattice to determine the optimal exercise strategies. Such algorithms are processed backward in time from the latest epoch to the earliest epoch, and such algorithms need to evaluate conditions occurring only at all states in the lattice, not along all paths through the lattice. From epoch 0 to epoch H, there are 2^H paths through a binomial lattice, but only $(H + 1)(H + 2)/2$ total states, if the lattice is connected. Thus, many option-pricing problems can be solved efficiently and accurately on a connected lattice. Without a lattice (whether connected or not), backward induction is not possible. From a purely mathematical viewpoint, it is difficult to construct *optimal* exercise strategies for many option problems by doing calculations on interest rate paths sampled from a continuous model. From a practical viewpoint, note that real-world options are exercised by people who manage portfolios or trading positions, or who run corporations or other businesses. The *behavior* of these people, as to their strategies for rational (if not mathematically optimal) exercise of the options they hold, can be modeled sufficiently accurately that the options are valued properly by way of calculations performed on paths sampled from a continuous model.

Many important problems involve *path-dependent* cash flows, for example, the pricing of prepayable mortgages and instruments derived from them, and the valuation of interest-sensitive insurance liabilities. For such problems, the possible paths of interest rates must be considered, not merely the possible states of the world. A connected lattice offers no special computational advantages in these situations. In fact, when path-dependent cash flows are involved and a lattice model is used, P paths of interest rates will have to be sampled, just as if a continuous model were being used.

In summary, several compelling factors favor the use of continuous-state models over discrete-state models: (i) a discrete model's lack of computational advantage in the common case of path-dependent cash flows; (ii) the need to use the same model consistently for all assets and liabilities, whether their cash flows are path independent or path dependent; and (iii) a continuous model's ability to sample states of the world sufficiently densely and to accommodate realistic yield curve dynamics.

4. Riskless Arbitrage, Complete Markets, and the Risk-Neutral World

4.1 Basic Concepts

This section is shorter than it could be, so that we can proceed to the main subject of the paper. The underlying mathematics are generally presented in an imposing manner and have been the subject of numerous lengthy seminal papers on the application of stochastic process theory to financial economics. The topic of riskless arbitrage is dealt with well in the paper by Pedersen, Shiu, and Thorlacius [18], and brief comments on the role of the risk-neutral world in option-pricing calculations can be found in the text by Cox and Rubinstein [6].

The concept of a *riskless arbitrage opportunity* is not difficult. If one asset or portfolio of assets can be sold and the proceeds of the sale can be used to purchase a different asset or portfolio of assets whose performance will be superior to that of the original asset or portfolio over a specified holding period (infinitesimal or finite depending on the situation), regardless of the states of the world during and at the end of the holding period, then a riskless arbitrage opportunity is said to exist. One need merely sell the first asset or portfolio and purchase the second to be guaranteed of having more wealth at the end of the holding period without having incurred greater risk. The reason that such an opportunity is said to be *riskless is* that wealth can be created without investing any capital at all by selling the first asset *short* (that is, selling it before purchasing it), and using the proceeds of the short sale to purchase the second asset. In this situation, there is no net outlay of funds, but there is a guarantee of positive wealth at the end of the holding period, because the second asset can then be sold for more than is then needed to cover (close out) the short position by purchasing the first asset.

Financial economists and other reasonable people assume that no riskless arbitrage opportunities exist in an equilibrium world. In other words, prices of assets are assumed to adjust continuously to eliminate opportunities for riskless arbitrage. For this to occur, a number of assumptions must be made: assets are perfectly divisible, unlimited short sales are possible, trading takes place continuously without transaction costs, investors act rationally and prefer more wealth to less wealth, and there are no taxes. Although these assump-

tions are quite stringent, one should not debate too strenuously whether small arbitrages can exist in the real world for brief periods because the assumptions are only approximations to reality. Instead, one should regard the concept of an equilibrium world in which riskless arbitrage opportunities do not exist as fundamental to constructing a sound financial theory for pricing assets.

To see how the concept of riskless arbitrage can lead to a theory for establishing the *relative* prices of assets, consider again the situation described above, modified slightly. Suppose that an asset for which one wants to establish the arbitrage-free price is equivalent to a portfolio of different assets for which one knows the prices. *Equivalence* is used in the sense that the performance of the single asset and that of the portfolio are *identical* over a specified holding period. Then it follows that the single asset and the portfolio of assets must have the same prices, or else there would be a riskless arbitrage opportunity, wherein the more expensive one could be sold short and the less expensive one purchased, guaranteeing a profit without taking any risk. Thus, one establishes the arbitrage-free price of the single asset as equal to the *known* price of the portfolio of assets. For this approach to be generally applicable and therefore lead to a pricing theory, it is necessary to assume that the financial markets are complete, meaning that any given asset is equivalent to some portfolio of fundamental assets.[1] This *replicating portfolio* might not be equivalent to the given asset over all holding periods. The portfolio's holdings might have to be adjusted from time to time, perhaps continuously, to maintain the equivalence. Having to rebalance the replicating portfolio is of no consequence, however, because the ability to trade continuously absent transaction costs, as assumed earlier, enables equivalence to be maintained without having to inject additional money into the portfolio; the replicating strategy is said to be *self-financing*.

If the financial markets are complete and no opportunities for riskless arbitrage exist, then the prices of all assets can be determined relative to the prices of their replicating portfolios. Under these assumptions, the relative prices of assets cannot depend on individual investors' *preferences,* which include their differing subjective views on the probabilities of occurrence of various future states of the world and their differing degrees of aversion to risk. Otherwise, riskless arbitrage opportunities would arise. Because relative asset prices must be *preference-free,* we can choose a frame of

reference in which the pricing of assets is particularly straightforward, namely, the *risk-neutral world*. It does not mean that we must adopt such a setting, only that we are permitted to do so, and that we will obtain the correct relative prices for assets if we do. Black and Scholes [3] derived their now-famous formula for the price of a call option on a share of non-dividend-paying stock in terms of the price of the underlying stock by applying the no-riskless-arbitrage condition to a combined position of buying the call option and selling short its replicating portfolio. They solved the resulting differential equation for the price of the call option after establishing appropriate boundary conditions. Only later did others show that a simpler derivation is possible by moving into the risk-neutral world and performing the pricing calculation there (for example, refer to [5]).

What is the risk-neutral world, and why are pricing calculations simpler there? In the risk-neutral world, investors do not require a premium for assuming risk. Thus, assets are priced at their expected present values. In other words, risk-neutral investors behave like traditional actuaries. When pricing assets, they project cash flows along interest rate paths, then discount the cash flows at the one-period interest rates occurring along the paths, and finally calculate the *expected present value* by weighting the present value for each path by that path's probability of occurrence.[2] Moreover, in the risk-neutral world, the probabilities of occurrence of the various paths do not depend on investors' subjective views of the likelihood that different future states of the world will arise. I now show how this description of the risk-neutral world can be used to construct an arbitrage-free model of interest rates.

In a binomial model, the states of the world represented at the nodes of the lattice can be determined from the assumed stochastic process for the one-period interest rate, for example, a discrete geometric Brownian motion. Then, the *risk-neutral probabilities* of up and down transitions at each node can be established to ensure that all zero-coupon bonds are priced properly by the expected-present-value algorithm described above. Alternatively, the no-riskless-arbitrage conditions can be used to establish the possible states of the world, given assumed risk-neutral probabilities for up and down transitions at each node; for example, $\theta_{i,t}$ for the up transition and $1 - \theta_{i,t}$ for the down transition at the i-th node at epoch t, with $0 < \theta_{i,t} < 1$. This is the approach used by Pedersen, Shiu, and Thorlacius [18]. In a continuous model, it is convenient to adopt the

approach of assuming that the risk-neutral probabilities are given, and then generating a *finite* number of interest rate paths appropriately. It is usual to generate *equal-probability* paths of interest rates by randomly sampling, epoch to epoch, from an assumed stochastic process, and to adjust, epoch by epoch, the distribution of interest rates to ensure that the no-riskless- arbitrage conditions hold.

4.2 Example: A One-Factor Lognormal Model of Short-Term Interest Rates

The rest of this section is devoted to an example in which the natural logarithm of the ratio of the one-period rate of interest at epoch t to the one-period rate of interest at epoch $t - 1$ is normally distributed with mean μ and standard deviation σ. It is conventional to refer to this example as a lognormal stochastic process for the one-period (spot) rate. The initial term structure of interest rates (all the spot rates or all the spot prices) is assumed to be specified exogenously. The objective is to generate an arbitrage-free set of P equal-probability paths of one-period interest rates out to epoch H, the assumed horizon for the desired application. In practical applications, limitations on computer memory and execution time usually constrain the choice of P to between 100 and 1000.

A single path of one-period interest rates can be created by starting from the given initial one-period rate, then randomly sampling from the assumed lognormal distribution a one-period rate at epoch 1 and using it as the starting one-period rate for randomly sampling from the assumed lognormal distribution a one-period rate at epoch 2, and so on, out to epoch H. Independently repeating this entire sequence of computations P times gives rise to P equal-probability paths of one-period interest rates. Unfortunately, the set of paths is *not* arbitrage free. To obtain an arbitrage-free set of paths, all P one-period rates at each epoch must be multiplied by an appropriate adjustment factor that is the same for all P rates, but that differs from epoch to epoch.[3] The proper approach involves generating and adjusting the one-period rates at epoch 1, which evolve from the given initial one-period rate at epoch 0; then generating and adjusting the one-period rates at epoch 2, which evolve from the adjusted one-period rates at epoch 1; and so on; and finally generating and adjusting the one-period rates at epoch H, which evolve from the adjusted

one-period rates at epoch $H - 1$. The arbitrage-free algorithm is described more precisely as follows:

Step 1. Using the initial (epoch 0) one-period rate and the assumed lognormal process, generate P random one-period rates at epoch 1. Call these the *epoch 1 unadjusted one-period rates*. Calculate P *epoch 1 adjusted one-period rates* by multiplying all P epoch 1 unadjusted one-period rates by the adjustment factor exp $[\lambda_1 - \mu]$, where λ_1 is determined so that the following condition holds: The expected present value of a $1 payment at epoch 2 is equal to the *given* initial spot price of the two-period zero-coupon bond. The expected present value is measured at epoch 0 and is based on the initial one-period rate and the epoch 1 *adjusted* one-period rates.

Step 2. Using the epoch 1 adjusted one-period rates and the assumed lognormal process, extend the P paths of interest rates one period by generating P random one-period rates at epoch 2. Call these the *epoch 2 unadjusted one-period rates*. Calculate the P *epoch 2 adjusted one-period rates* by multiplying all P epoch 2 unadjusted one-period rates by the adjustment factor exp $[\lambda_2 - \mu]$, where λ_2 is determined so that the following condition holds: The expected present value of a $1 payment at epoch 3 is equal to the *given* initial spot price of the three-period zero-coupon bond. The expected present value is measured at epoch 0 and is based on the initial one-period rate and the epoch t *adjusted* one-period rates for $t = 1, 2$.

....

Step H. Using the epoch $(H - 1)$ adjusted one-period rates and the assumed lognormal process, extend the P paths of interest rates one period by generating P random one-period rates at epoch H. Call these the *epoch H unadjusted one-period rates*. Calculate the P *epoch H adjusted one-period rates* by multiplying all P epoch H unadjusted one-period rates by the adjustment factor exp $[\lambda_H - \mu]$, where λ_H is determined so that the following condition holds: The expected present value of a $1 payment at epoch $H + 1$ is equal to the *given* initial spot price of the $(H + 1)$-period zero-coupon bond. The expected present value is measured at epoch 0 and is based on the initial one-period rate and the epoch t *adjusted* one-period rates for $t = 1, 2, ..., H$.

The expected present values referred to in steps 1 through H above are calculated as described earlier in this section: that is, as the simple arithmetic mean (since all paths have equal probabilities $1/P$) of the P present values obtained by discounting the cash flows occurring along each path at the relevant one-period interest rates occurring along the path. The sequence λ_1,

$\lambda_2, ..., \lambda_H$ is referred to as the *time-dependent drift* of the stochastic process. Thus, constraining the process to be arbitrage free is tantamount to setting its time-dependent drift so that all zero-coupon bonds are priced at values equal to those derived from the exogenously specified initial yield curve. Otherwise, trivial riskless arbitrages would exist. Note that all traces of the originally assumed drift μ are eliminated by the adjustment factors exp $[\lambda_t - \mu]$. In essence, the drift is reset from μ to the appropriate λ_t at each epoch t to ensure that the model is arbitrage free.

Performing steps 1 through H ensures that the expected-present-value algorithm will produce the observed (exogenous) initial prices for all zero-coupon bonds having maturities less than or equal to $H + 1$ periods. As a consequence, any stream of fixed and certain cash flows occurring at epochs up to $H + 1$ will be priced fairly relative to the zero-coupon bonds, because an arbitrary stream of fixed and certain cash flows is equivalent to some portfolio of zero-coupon bonds. Similarly, the arbitrage-free set of paths of one-period interest rates and the expected-present-value algorithm can be used to price any interest-rate-contingent stream of cash flows occurring at epochs up to H. (In general, one cannot get as far as epoch $H + 1$, because the cash flows at epoch $H + 1$ may depend on the adjusted one-period rates at epoch $H + 1$, and those rates have not been determined.) The prices so obtained for interest-rate-contingent cash flow streams will be arbitrage free relative to the prices of the zero-coupon bonds, because any such stream can be shown to be equivalent to some dynamically adjusted replicating portfolio consisting only of zero-coupon bonds.

One has to be careful about what is meant by the descriptor "arbitrage free." The H-step algorithm ensures a fair price at *epoch 0* for any stream of cash flows not extending beyond epoch H. The epoch 0 price of the cash-flow stream is said to be "fair" relative to the epoch 0 prices of all zero-coupon bonds, because opportunities for riskless arbitrage between the cash-flow stream and the set of zero-coupon bonds have been eliminated through the application of the H-step algorithm. The key question, however, is whether the H-step algorithm ensures arbitrage-free dynamics for the *full* yield curve at all epochs or just for the one-period rate at all epochs. The answer can be either yes or no, depending on the objective. If some control over the statistical properties of yield curve movements is desired, such as being able to specify exogenously the

correlation between the movements of the one-period rate and the current-coupon yields at all other maturities, then the answer is no: the H-step algorithm will not produce interest-rate behavior consistent with the externally supplied assumptions. The reason is that the stochastic behavior depends on only *one factor:* the dynamics of the one-period rate determines the dynamics of all other parts of the yield curve.

To see this, suppose we are situated at an arbitrary state of the world (namely, an arbitrary value of the one-period rate) at epoch t, where $t < H$. We determine the prices, and hence the yields, of all zero-coupon bonds with maturities up to $H + 1 - t$ periods by generating a large number of paths of one-period rates *emanating from this state* and then using the expected-present-value algorithm. In generating the paths emanating from the arbitrary state, we *are forced to use* the drift parameters $\lambda_1 + 1$, ..., λ_H that have *already* been determined. Thus, utilizing the H-step algorithm establishes the dynamics of the full yield curve in this one-factor model of interest rates. The model is fully arbitrage free at the initial state and at all future states of the world, but a huge sacrifice is necessary to achieve this result: the ability to specify the desired yield curve dynamics must be (at least partially) surrendered!

Before generalizing the one-factor model to a multiple-factor model and therefore gaining control over yield curve behavior, we consider how the H-step algorithm described above achieves arbitrage-free stochastic interest-rate dynamics. Earlier, we stated that a model of interest rates can be rendered arbitrage free by either (i) assuming the states of the world to be given and the risk-neutral probabilities to be determined, or (ii) assuming the risk-neutral probabilities to be given and a finite number of states of the world to be determined appropriately. The H-step algorithm uses the latter approach. It assumes that the P paths have equal probability, and then determines P one-period rates at each epoch in a manner to ensure that the no-riskless-arbitrage conditions are satisfied. Because theoreticians generally speak and write in terms of the former approach, it makes sense to ask whether there is some relationship between the approaches. The answer is yes. Take the limit as P tends to infinity and ensure that all possible paths of one-period interest rates are represented. In this case, because the full continuous distribution of states at each epoch is represented, the H-step algorithm actually establishes the probability distribution of paths in the risk-neutral world. Thus, when all possible paths are represented, solving for the risk-neutral drift of the stochastic process is equivalent to establishing the risk-neutral probabilities.

5. Constructing an Interest Rate Generator

This section extends the approach described in Section 4.2 to allow different parts of the yield curve to be modeled simultaneously. This is important in many practical applications. For example, prepayments arising from a seasoned pool of fixed-rate residential mortgages are commonly modeled as depending on the currently prevailing and recent levels of the ten-year U.S. Treasury bond yield. To project the cash flows from the mortgage pool, a stochastic model of ten-year U.S. Treasury bond yields is needed. To price the pool of mortgages properly, that model must be arbitrage free. Because movements of short-term U.S. Treasury bill yields and ten-year U.S. Treasury bond yields are correlated, the models of the short-term yields and ten-year yields must be developed together and consistently. Other similar examples abound: insurance company interest-sensitive life and annuity products are often priced relative to prevailing intermediate-term U.S. Treasury note yields; most adjustable-rate preferred stocks have quarterly dividends that reset off the highest of the three-month U.S. Treasury bill yield, the ten-year U.S. Treasury note yield, and the 30-year U.S. Treasury bond yield; and many options are written on intermediate-term U.S. Treasury notes and bonds.

Even when many parts of the yield curve are modeled together, the one-period rates continue to play a special role. In the expected-present-value algorithm, only the one-period rates are used to discount cash flows. Yields of bonds with maturities greater than one period affect the arbitrage-free price of a stream of interest-rate-contingent cash flows only to the extent that the amounts of the cash flows depend on those yields. The examples listed above illustrate such dependence.

5.1 What Random Variables Should Be Modeled?

Before describing *how* to build a stochastic interest rate generator, we must decide *what* interest rates are to be modeled. From Section 2, the choices are: current-coupon yields, spot rates, or forward rates. Because

these three sets of rates are equivalent, any one set may seem to be as good as another for stochastic modeling purposes, but that is not the case. From the recent literature, it is clear that academicians prefer to model forward rates. It is useful to understand why.

If one views interest rates as the *cost of money*, it is reasonable to constrain them to be non-negative. Some models permit negative interest rates (Vasicek [23] and Ho and Lee [11], for example), but most theoreticians and practitioners regard non-negativity of interest rates as a basic constraint that an acceptable model must satisfy. If all one-period $(n,1)$ forward rates, as defined in Section 2, are non-negative, then all spot rates and all forward rates are non-negative. Thus, modeling forward rates by a process that ensures that they are non-negative also ensures that spot rates and current-coupon yields are non-negative. Modeling spot rates by a process that ensures that they are non-negative does *not* ensure that all forward rates will be non-negative. For example, if both the $(n + 1)$-year and the n-year spot rates are positive, but the n-year spot price is less than the $(n + 1)$-year spot price, then the $(n,1)$ year forward rate will be negative, as discussed in Section 2. Modeling current-coupon yields by a process that ensures that they are non-negative can produce a result even more *pathological*, because negative spot prices can occur, leading to spot rates that are *complex numbers* (having both *real* and *imaginary* parts).

The preceding points would seem to make a clear case for modeling forward rates, and that is the route that authors of most recent papers have chosen. (Heath, Jarrow, and Morton [8] is an example.) From a practical viewpoint, the trouble with modeling forward rates is that few people seem to have an intuitive feel for how forward rates should behave. Traders in U.S. government securities often do not know what a theoretician means by a one-period forward rate, and they certainly do not have an intuitive feel for the relative volatilities of the various $(n,1)$ period forward rates. Nor do they have a view on whether the movements of the $(n,1)$ period forward rates for different n are weakly or strongly correlated. These observations cannot be ignored if the goal is to build a stochastic generator that will be useful in practical situations.

What fixed-income traders and portfolio managers understand well is the *yield curve*. Thus, it seems reasonable to build a model using current-coupon yields as the random variables. Of course, the architects and engineers of such models have a strong duty to test how

frequently pathological spot rates and forward rates arise when reasonable ranges of input assumptions are used, and to reject models for which that frequency is too high (greater than 1 percent, for example). A rejection criterion is necessary, because the so-called pathologies are actually just different forms of riskless arbitrage opportunities which we have so assiduously eliminated elsewhere.

A model that allows the user to input assumptions about the behavior of current-coupon yields, but which, perhaps unknown to the user of the model, simulates forward rates of interest as the random variables, is the ideal solution to the dilemma posed above. This is undoubtedly the proper course to follow, but I have not yet discovered a wholly satisfactory method for translating input assumptions relating to current-coupon yield volatilities, correlations, and mean reversion strengths (all defined in Section 5.2) into model parameters that apply to a specified stochastic process followed by forward rates. Because this is likely to be the natural evolution of interest rate models, I should remark that the stochastic process introduced in Section 5.2 is as valid for spot rates or forward rates of interest as it is for current-coupon yields. Thus, most of the material presented in Sections 5.2 and 5.3 can be read as if it applies directly (or with obvious modifications) to spot or forward rates of interest.

5.2 A Continuous Stochastic Model of Current-Coupon Yields

Let $r_{k,t}$ denote the random variable for the yield at epoch t of the non-callable current-coupon bond that has a maturity of k periods.[4] The natural logarithm of $r_{k,t}$ is assumed to obey the following process:

$$\ln r_{k,t+1} = \lambda_{k,t+1} + (1 - \phi_k)\ln r_{k,t} + \sigma_k \varepsilon_{k,t+1}$$

where $\lambda_{k,t+1}$ is the drift that was introduced in Section 4, ϕ_k measures the strength of mean reversion ($0 \le \phi_k \le 1$), σ_k is the logarithmic yield volatility, and $\varepsilon_{k,t+1}$ is a random standard normal deviate (that is, a random normal deviate with zero mean and unit variance). It is assumed that $\varepsilon_{j,s}$ and $\varepsilon_{k,t}$ are independent random variables if $s \ne t$ and that they have linear correlation coefficient ρ_{jk} if $s = t$, with $\rho_{jj} = 1$.

The subscript k for the random yield variables and the parameters has been used to indicate that there is one equation for each current-coupon bond maturity. Each yield r_k follows its lognormal random walk with

its own mean reversion ϕ_k, volatility σ_k, and time-dependent drift $\lambda_{k,t}$. Using a lognormal process ensures that no r_k can ever be negative. The stochastic processes followed by the yields of the current-coupon bonds of the various maturities are *not* independent of each other; the contemporaneous shocks ε_j and ε_k to the natural logarithms of the current-coupon yields r_j and r_k, respectively, have correlation ρ_{jk}.

That ϕ_k can be interpreted as the strength of mean reversion for the yield r_k becomes a little clearer if the drift is time-independent and given by $\lambda_{k,t} = \phi_k \ln \mu_k$ at all epochs t, where μ_k is the *mean reversion level* of r_k, apart from a factor that is generally close to unity if ϕ_k is not too small. An explicit expression for the mean reversion level is given in Section 5.3. The larger is ϕ_k, the faster the expected value of r_k approaches its long-run limit. Thus, ϕ_k measures the strength of the mean reversion. When $\phi_k = 0$, the yield r_k follows a pure lognormal random walk with drift, and there is no finite long-run mean level of yields. When $\phi_k = 1$, the yield r_k has no memory whatsoever of its levels at previous epochs and, apart from time-dependent drift, has the same lognormal probability distribution at each epoch.

The formula for $r_{k,t+1}$ is a *first-order* difference equation for which the initial condition is the initial current-coupon yield $r_{k,0}$. The difference equation is said to be *first-order autoregressive* because the value of $r_{k,t+1}$ can be determined, apart from the random shock $\varepsilon_{k,t+1}$, by looking back only one period to the value of $r_{k,t}$. In practical applications, the difference equation is used as the basis for iterative random simulation. A single path of interest rates can be generated randomly by starting from $r_{k,0}$ and iterating one period at a time to obtain $r_{k,1}$, $r_{k,2}$, ..., $r_{k,H}$. The set of $r_{k,t}$ for all k defines the yield curve at epoch t, and the sequence of yield curves at epochs 0, 1, 2, ..., H defines the interest-rate path. The stochastic nature of the interest-rate path derives from the random shocks ε that occur from epoch to epoch along the path.

Consider a single path of interest rates. In generating the path on a computer, a column vector ε of standard normal deviates must be generated randomly at each epoch. There is one element of ε for each maturity index k. If the correlation matrix ρ is positive definite, it can be factored as $\rho = LL^T$, where L is a lower triangular matrix. This is known as the Choleski or square-root factorization of ρ (refer to [4]). Suppose ξ represents a column vector of *independent* random standard normal deviates. Then, the column vector $\varepsilon = L\xi$ of random standard normal deviates has a correlation matrix equal

to the given matrix ρ, as desired. The column vector ξ can be generated by any acceptable method for producing random standard normal deviates, such as the polar method [24].

The stochastic process specified above possesses a useful *scaling* property that allows the interest-rate model to be adjusted easily to whatever time interval between epochs is desired. If the difference equation is applied s *times* rather than once, it takes the form:

$$\ln r_{k,t+s} = \Lambda_{k,t+s} + (1 - \Phi_k)\ln r_{k,t} + \phi_k \Sigma_k E_{k,t+s}$$

where $\Lambda_{k,t+s}$ depends on $\lambda_{k,u}$ for $1 \le u \le s$, ϕ_k, and s; Φ_k depends on ϕ_k and s; Σ_k depends on σ_k, ϕ_k, and s; and P_{jk}, the scaled linear correlation coefficient, depends on ρ_{jk}, ϕ_j, ϕ_k, and s. $E_{k,t+s}$ is a random standard normal deviate. Thus, under a scaling of the time interval from 1 to s, the same form of process obtains, but its parameters must be scaled appropriately relative to the parameters for the unscaled process. To emphasize the similarity, I have used upper-case Greek letters for the scaled process to correspond to the lower-case Greek letters for the unscaled process: specifically, $\lambda \leftrightarrow \Lambda$, $\phi \leftrightarrow \Phi$, $\sigma \leftrightarrow \Sigma$, $\rho \leftrightarrow P$, and $\varepsilon \leftrightarrow E$.

What parts of the yield curve should be modeled? If a quarter-year time interval were used, for example, one could model the yields for all current-coupon bonds with maturities at quarter-year intervals up to the maximum maturity bond needed for the particular application. In fact, it might seem that the yields for all bonds *must* be modeled if a fully arbitrage-free set of paths is desired. This is true if one insists that each random yield strictly follow the specified stochastic process. Modeling all maturities leads to a very large computing problem, as to both memory requirement and execution time. I believe that it is preferable to model only what Ho [10] has called *key (spot) rates* or *key (current-coupon) yields* and to obtain all other needed rates or yields by linear interpolation. Reitano describes the same approach in a recent paper [20].

The justification for not using all bond maturities in the model is that fixed-income traders generally use the following new-issue or outstanding U.S. Treasury securities as *benchmarks:* bills having maturities of 3, 6, and 12 months; notes having maturities of 2, 3, 5, 7, and 10 years; and bonds having maturities of 20 and 30 years. These are the *key maturities* that should be modeled. The reference yields for hypothetical current-coupon instruments with other maturities are usually quoted by traders on the basis of linear interpolation between the

key yields. A result of using the key-yield-with-linear-interpolation approach is that the random yields of bonds with other than key maturities will only approximately follow the form of stochastic process that is precisely followed by the key yields, and the non-key yields will not be arbitrage free. This sacrifice, made in response to the computational limitations of most computers, will likely be acceptable only if sufficiently many key maturities are used and if those maturities are spaced sufficiently closely.

5.3 Adjusting the Drift Parameters

How to set the drift parameters depends on the application. For example, to permit investors to base their analyses of the risk-reward tradeoffs among various risky investment strategies on their own subjective views regarding the behavior of interest rates, the model can be tuned to characterize the real world. This is accomplished by setting the drift parameters appropriately. A commonly used assumption in the real world is that the drift is time independent: specifically, that $\lambda_{k,t} = \phi_k \ln \mu_k$ at all epochs t. In this case, by exponentiating the difference equation for $\ln r_{k,t+s}$, then taking the mathematical expectation, and finally taking the limit as s tends to infinity, it can be seen that the long-run means of the random variables r are equal to $\mu\exp[\sigma^2/2(1 - (1 - \phi)^2)]$, where the subscripts k have been omitted for notational convenience. These are the mean reversion levels of the key yields. Alternatively, the investor can set a *time-dependent* drift of the process to reflect his or her short-term and intermediate-term views of the trend of the expected levels of the key yields, followed by a *time-independent* drift to establish the desired long-term mean reversion levels of the key yields.

The term "real world" is not intended to convey either the notion of *always* being the "proper" or "superior" framework for analysis or the notion that other worlds are somehow *always* "improper" or "inferior" frameworks for analysis. The real world is an environment in which different investors have differing subjective views, differing degrees of risk aversion, and differing utility functions. Thus, the real world differs from the risk-neutral world described in Section 4, which is an *artificial* construct that enables fair prices to be calculated for assets and liabilities (whether they are potentially "risky" or not) in the simplest manner possible: specifically, by way of the expected-present-value algorithm. The analysis of risk-reward tradeoffs among

various risky strategies is carried out in the real world by (i) constructing a *frontier* of strategy possibilities, all of which are *efficient* in the sense of having the least risk for a given expected return, and then (ii) selecting the optimal efficient strategy on the basis of one's own utility function. These important concepts of modern portfolio theory are covered in the text by Sharpe and Alexander [22].

If the intended use of the interest rate model is to determine the arbitrage-free price of an interest-rate-contingent stream of cash flows, the drift parameters must be set appropriately to characterize the risk-neutral world. Adjusting the drift parameters properly is a more difficult problem when many key yields are involved than when only the one-period rate is involved. In order to obtain yield curve dynamics that exhibit sufficiently many degrees of freedom (by means of correlated stochastic processes for the key maturities), and yet still ensure that the model is arbitrage free (at the key maturities), the drift parameters must be *state dependent,* not merely time dependent. The principal objective in adjusting the drift parameters of the processes is to ensure that the expected-present-value algorithm prices correctly the exogenously specified *initial yield curve.* However, *any future state* must also be viewed, at the epoch it occurs, as an "initial" state, and the drift parameters must be adjusted properly from that state forward in time in order to ensure that the expected-present-value algorithm also prices correctly the yield curve at the new "initial" state.

The arbitrage-free algorithm is completely specified by showing how to start from an arbitrary "initial" state and create fully arbitrage-free yield curves *one period later.* The local drift of each of the correlated processes is adjusted to satisfy the no-riskless-arbitrage condition of *step 1* in the method described in the next paragraph. (Only *step 1 is* ever used in this approach, but it is used at the true initial state and at each future state occurring at each future epoch out to the horizon.) Fortunately, the procedure does *not* require that a sample of paths be generated from each "initial" state. Instead, the drift adjustment is accomplished by computing the relevant expected present values in the no-riskless-arbitrage conditions as *integrals* over the *entire* continuous conditional lognormal distribution of states *one period ahead.* The integrals can be evaluated easily and accurately by the Gauss-Hermite method [9], and the local drift parameters can be determined efficiently by Newton-Raphson iteration [9]. If spot rates of interest are

assumed to be locally lognormally distributed, the computations are practical on a mainframe computer or its equivalent. However, when either current-coupon yields or forward rates of interest are assumed to be locally lognormally distributed, the number of independent stochastic factors in the model must be limited to a small number, such as two or three, for the computations to be feasible. Note that bonds of all maturities can be studied in a two-factor or three-factor model, but the "degrees of freedom" embedded in the variance-covariance matrix that is constructed from the volatility vector σ, and the correlation matrix p will be exactly two or three, respectively. Further details of the iterative method for obtaining a set of paths sampled from a model that is fully arbitrage free at the initial state and all future states of the world will be provided in a subsequent technical note. The essential points are covered in an article by Miller [17].

A yield curve model that is commonly used because it is computationally less demanding to construct than the model described above, but which is not free of riskless arbitrage opportunities at future states of the world, can be obtained by generalizing the no-riskless-arbitrage conditions discussed in Section 4. The objective is to generate a set of P equal-probability paths of yield curves (defined at the key maturities k) out to epoch H, the assumed horizon. The procedure is essentially the H-step inductive algorithm discussed in Section 4. The *exogenously given* initial set of key yields $r_{k,o}$ is the starting point. Step number t in the procedure assumes that P adjusted yield curves have been obtained at all epochs from 0 to $t - 1$. Using (i) the P adjusted yield curves at epoch $t - 1$, (ii) the assumed stochastic process, and (iii) trial values of the drift parameters $\lambda_{k,t}$ for all key maturities k, the interest rate paths are extended one period by generating P unadjusted yield curves at epoch t. Adjusted yield curves at epoch t are obtained by solving for the values of the drift parameters $\lambda_{k,t}$ that force the following no-riskless-arbitrage condition to hold for *all key maturities* k: At epoch 0, the price of the $(k + t)$-period zero-coupon bond, as computed by applying the expected-present-Value algorithm, must equal the *given* initial price of the $(k + t)$-period zero-coupon bond, that is, the $(k + t)$-period spot price that is derived from the *given* initial yield curve. The following technical points are intended to be helpful to anyone who attempts to program the algorithm:

• To apply the no-riskless-arbitrage condition, the *spot-prices* at key maturities must be calculated. This means that yield curves specified at all maturities, not just key maturities, must be obtained before they can be transformed into the corresponding curves of spot rates and spot prices. For any maturity that is between two adjacent key maturities, the associated current-coupon yield is found by interpolating linearly between the relevant key yields, as discussed earlier.

• The most difficult part of implementing the algorithm is solving for the values of the drift parameters $\lambda_{k,t}$ that force the no-riskless-arbitrage conditions to be satisfied. The method of false position, also known as *regula falsi* [9], has been found to give satisfactory results in that it succeeds in finding the roots and usually converges rapidly.

• It was pointed out in Section 5.1 that constraining current-coupon yields to be non-negative does not preclude negative spot prices from arising. If they do occur, trouble is likely to be encountered when solving for the drift parameters. In particular, difficulties arise in the handling of complex numbers. Remedies can be devised for these annoyances, but a description of them is beyond the scope of this paper.[5]

5.4 Variance-Reduction Techniques

Because the model described in Section 5.2 must be implemented on a computer, either memory capacity or execution time will limit the number of interest rate paths that can be generated for a given application. Fortunately, calculating the arbitrage-free price of an interest-rate-contingent cash-flow stream is an *expected value* problem. Only the *mean of* the distribution of present values across paths needs to be determined. None of the higher moments of that distribution matter. Nevertheless, the *entire* probability distribution of interest-rate paths, not just its mean, affects the arbitrage-free price of a general cash-flow stream. Thus, it is important to sample interest-rate paths efficiently, and standard variance-reduction techniques are useful in this regard. A good reference for this subject is the text by Rubinstein [21].

The methods of *antithetic variates* and *stratified sampling* were tested on the simple problem of estimating the moments of the probability distribution of the sum of 40 independent random standard normal deviates and then on the more realistic problem of pricing European bond options. Each variance-reduction technique was applied separately, and then they were applied together. All three cases were compared against the control experiment of

crude Monte Carlo sampling, that is, pure random sampling without any variance reduction technique applied. In all situations, an even number of paths $P = 2Q$ was used, where Q is a positive integer greater than one. The algorithms that were used in selecting a random sample of size P are described below for the three variance-reduction methods that were studied.

Antithetic Variates

A sample of Q independent standard normal deviates $(n_1, n_2, ..., n_Q)$ is chosen randomly. Another set of Q standard normal deviates is obtained by reversing the signs of the first Q deviates. (This step is what gives the technique its name.) Together, the two sets of Q deviates form the required random sample of P deviates, which can be displayed descriptively in *pairs* as: $(n_1, -n_1; n_2, -n_2; ...; n_Q, -n_Q)$. This technique results in *mirror-image* pairs of paths of the natural logarithms of the key yields.

Stratified Sampling

The positive half of the standard normal density is divided into Q strata of equal probability $1/P$. A stratified sample $(n_1, n_2, ..., n_Q)$ is created by placing n_i for $i = 1, 2, ..., Q - 1$ at the midpoints of the first $Q - 1$ strata, and n_Q at such point in the Q-th stratum as to ensure unit variance for the full sample of P deviates. The required random sample of P deviates is obtained by *shuffling* the ordered stratified sample $(-n_Q, ..., -n_2, -n_1, n_2, ..., n_Q)$ randomly. Any trace of nonrandom pairing disappears completely after random shuffling.

Antithetic Variates Combined with Stratified Sampling

A stratified sample of Q positive standard normal deviates is created as previously described, and each deviate in that sample is *separately and independently* multiplied by +1 with probability 1/2 or by −1 with probability 1/2. The resulting sample is then *shuffled* randomly to produce the sample $(n_1, n_2, ..., n_Q)$, which is augmented by another set of Q deviates that is obtained by reversing the signs of all the n_i for $i = 1, 2, ..., Q$. The required random sample of P deviates is displayed descriptively in *pairs* as: $(n_1, -n_1; n_2, -n_2; ...; n_Q, -n_Q)$. This technique results in *mirror-image* pairs of paths, in the sense described above.

In the test situations in which paths of interest rates were used, *independent* samples of P random standard normal deviates were needed for each key maturity k at each epoch t, regardless of the sampling method studied. When stratified sampling was used, the stratified sample of Q positive standard normal deviates needed to be created only once, and that sample was used repeatedly in the second and third algorithms described above. There are several ways to combine the methods of antithetic variates and stratified sampling. The approach described above has both advantages and disadvantages relative to other algorithms for combining the techniques.

The effectiveness of the variance-reduction techniques was tested for sample sizes of $P = 100$ and $P = 1000$. Unfortunately, the results of the tests were ambiguous. Each of the three methods generally produced satisfactory results relative to using crude Monte Carlo sampling, but none of the approaches emerged as clearly the most effective. Due to the *pairing* of deviates that was described above, any method that involves antithetic variates always reduces the effective sample size to one-half of the sample size for either crude Monte Carlo or stratified sampling. Thus, when a given statistic is estimated on the basis of a fixed sample size, the standard deviation of the error that results when the method of antithetic variates is involved is generally larger, by a factor of $\sqrt{2}$, than the standard deviation of the error that results when either crude Monte Carlo or stratified sampling is used. However, the method of antithetic variates is designed to calculate some statistics *exactly*, regardless of sample size.[6] It is not always clear which of these opposing effects on the size of the estimation error will dominate. More analysis is needed before stronger conclusions can be made about which variance-reduction technique is to be preferred in a given situation.

6. Using an Interest Rate Generator

The first task that must be accomplished before the interest rate generator described in Section 5.2 can be used is to estimate the parameters of the model and to input other necessary assumptions. The generator requires as input an initial yield curve $r_{k,0}$ that is specified at the key maturity indexes k. Although the parameters of the model can be chosen arbitrarily, it is reasonable to begin by using values derived from a historical analysis of yield curve movements. Performing

linear regressions on at least the last two years of weekly yield curve data will provide estimates of the mean reversion levels μ and strengths ϕ, the yield volatilities σ, and the yield shock correlation matrix ρ, as were defined in Sections 5.2 and 5.3.[7] The scaling transformations described in Section 5.2 then must be applied to adjust the parameters estimated from the historical data to values appropriate to the time interval used in the stochastic difference equation, for example, from weekly periods in the historical data to quarter-year periods in the difference equation. When the partially arbitrage-free version of the yield curve model described in Section 5.3 is used, the no-riskless-arbitrage conditions override the mean reversion level parameters μ and result in values of the drift parameters λ that depend on the initial yield curve. When the fully arbitrage-free version of the yield curve model described in Section 5.3 is used, the no-riskless-arbitrage conditions override the mean reversion strength parameters ϕ as well, thus eliminating the explicit effects of mean reversion from the model.

Once the parameters and the initial yield curve have been specified, the model can be used to generate an arbitrage-free set of interest-rate paths. This set of paths and the expected-present-value algorithm defined in Section 4 are all that is needed to be able to project and price streams of interest-rate-contingent cash flows: callable corporate bonds, mortgage-backed securities, CMOs, floating-rate and indexed bonds, futures, options, interest rate swaps, caps, and floors, and all forms of liabilities. By perturbing or shocking the initial yield curve $r_{k,0}$ in different ways, one can compute various indexes that measure interest-rate risk. As an illustration, suppose that the arbitrage-free set of paths has been used to value a mortgage-backed security, and the resulting price is P. Now, suppose that the initial bond-equivalent current-coupon yield curve obtained by linearly interpolating the key yields $r_{k,0}$ is transformed into the equivalent term structure with spot rates expressed as forces of interest (that is, on a *continuously compounded* basis), and that these spot rates are all given the same small upward shock Δ—this is referred to as a *parallel* shock to the term structure. If key yields are taken from the yield curve that is equivalent to the parallel-shocked term structure, and are used as a *revised* initial yield curve from which a new arbitrage-free set of paths is created, a revised price P' can be calculated for the mortgage-backed security. The *effective duration* index D of the mortgage-backed security is given

by $D = (P-P')/P\Delta$.[8] Similarly, by using two different parallel shocks to the term structure, the *effective convexity* index of the mortgage-backed security can be calculated.

Risk indexes that measure price responses to fairly general *non-parallel* shifts in the term structure can be computed as described by Reitano ([19] and [20]). This is accomplished through shocking by amount Δ only the spot rate at the key maturity k, without shocking any of the spot rates at the other key maturities. Linearly interpolating the shocked term structure, transforming it back into a revised initial yield curve, and repricing the illustrative mortgage-backed security, gives a revised price P'_k. The *partial duration* or *key rate duration* D_k is given by $D_k = (P - P'_k)/P\Delta$. The sum of the partial durations D_k for all key maturities k *is* equal to D, the *total duration* index calculated above. A risk index that measures the sensitivity of the illustrative, mortgage-backed security's price to a small change in yield volatilities can also be calculated. This can be done for each key yield volatility σ_k separately to give *partial volatility durations* or for a constant shock to all key volatilities simultaneously to give a *total volatility duration*. The same approach can be applied to the yield curve shock correlation coefficients p_{jk} to calculate other risk indexes if desired.

Price and risk index calculations of the type described above have been utilized by Griffin [7] to develop an *excess spread* methodology for measuring profitability and its exposure to interest rate risk. In his paper, the new approach is applied to the case of an interest-sensitive annuity business. An interest rate generator of the type specified in Section 5.2, and adjusted to be arbitrage free as described in Section 5.3, is central to implementing Griffin's methodology, and thus is an essential tool in asset-liability management. To determine a financial institution's risk-neutral position and its current risk exposure relative to that risk-neutral position, the stochastic generator must be run in arbitrage-free mode by setting the drift parameters as described in Section 5.3. To analyze whether risky product-investment strategies should be adopted, the stochastic generator can then be run in real-world mode by adjusting the drift parameters to reflect management's views on expected interest rate behavior over time. Some institutions will be sufficiently risk averse that they will choose the risk-neutral position as their strategy. Or they may have insufficient capital to adopt prudently any strategy other than the risk-neutral position. Some institutions will have both the inclination

and the financial strength to adopt a risky strategy, and can use the generator in the real-world mode to analyze the relative attractiveness of different risky strategies.

7. Key Conclusions

This paper has shown that it is possible to gain a practical understanding of key concepts in financial economics without having to resort to a study of the mathematics of stochastic processes. Based on the assumptions that opportunities for riskless arbitrage do not exist and that the financial markets are complete, it was shown that a theory could be developed for pricing interest-rate-contingent cash-flow streams relative to the prices of all zero-coupon bonds, which are taken as exogenous inputs to the theory. The arbitrage-free prices are calculated in a straightforward manner by the expected-present-value algorithm, a technique that lies at the heart of actuarial science.

From a theoretical viewpoint, a model of interest rates should be based on forward rates of interest, rather than on spot rates of interest or current-coupon yields, because pathologies such as negative or complex interest rates are then eliminated automatically. From a practical viewpoint, however, model parameters are specified more naturally in terms of the statistical properties (volatilities and correlations) of current-coupon yields which serve as the reference points for traders who make markets in fixed-income securities. More research is needed in order to build a single model that combines the advantages of each approach. This is likely to be the next significant practical evolution of interest rate path generators.

Continuous-state models have important practical advantages over discrete-state models (and have no significant practical disadvantages): (i) realistic yield curve dynamics are accommodated more easily in a continuous model because a separate stochastic process equation is used for each key part of the term structure; (ii) many more states are sampled at the early epochs (and perhaps all the way out to the investor's horizon) in a continuous model, because paths of interest rates are the primary feature of the model; and (iii) moving between the risk-neutral world and the real world is accomplished more readily in a continuous model, because only the local drift of the process need be adjusted appropriately.

The drift of a stochastic process is an important "dial" in an interest-rate model that is based on the pro-

cess. The drift can be "tuned" to the risk-neutral world; so that application of the expected-present-value algorithm will result in arbitrage-free prices for interest-rate-contingent cash flow streams. This setting of the dial enables an asset-liability manager to determine the risk-neutral position properly. The drift can also be "tuned" to the real world, so that subjective views regarding the behavior of interest rates can be accommodated. This setting of the dial enables the asset-liability manager to analyze properly the risk-reward tradeoffs inherent in adopting risky positions.

The type of model of the yield curve presented in this paper can be used to price general streams of interest-rate-contingent cash flows and to compute indexes of interest rate risk, such as duration and convexity. The calculation of partial durations for the situation of non-parallel changes in the term structure is straightforward and natural under the model presented. Using the model involves generating a finite sample of interest rate paths on which to perform the price and risk index calculations. To improve the efficiency of the price estimation, variance-reduction techniques must be used when selecting the sample of paths. More research is needed to determine which techniques work best in given situations.

References

1. BALZER, L.A. "Term Structure of Volatility in Financial Markets," *Proceedings of the 1st AFIR Colloquium* 2 (1990): 57-63.
2. BLACK, F., DERMAN, E., AND TOY, W. "A One-Factor Model of Interest Rates and Its Application to Treasury Bond Options," *Financial Analysts Journal* 46, no. 1 (1990): 33-39.
3. BLACK, F., AND SCHOLES, M. "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy* 81 (1973): 637-59.
4. BURDEN, R.L., AND FAIRES, J.D. *Numerical Analysis*, 4th ed. Boston, Mass.: PWS Publishers, 1989.
5. COX, J.C., ROSS, S.A., AND RUBINSTEIN, M. "Option Pricing: A Simplified Approach," *Journal of Financial Economics* 7 (1979): 229-63.
6. COX, J.C., AND RUBINSTEIN, M. *Option Markets*. Englewood Cliffs, N.J.: Prentice Hall, 1985.
7. GRIFFIN, M.W. "An Excess Spread Approach to Nonparticipating Insurance Products," *TSA* XLII (1990): 229-46.

8. HEATH, D., JARROW, R., AND MORTON, A. "Bond Pricing and the Term Structure of Interest Rates: A New Methodology," Working Paper, Johnson Graduate School of Management, Cornell University, 1988.

9. HILDEBRAND, F.B. *Introduction to Numerical Analysis*, 2nd ed. New York, N.Y.: Dover Publications, Inc., 1987.

10. HO, T.S.Y. "Key Rate Durations: A Measure of Interest Rate Risk Exposure," Working Paper, Global Advanced Technology Corporation, 1990.

11. HO, T.S.Y., AND LEE, S.-B. "Term Structure Movements and Pricing Interest Rate Contingent Claims," *Journal of Finance* 41 (1986): 1011-29.

12. HUANG, C.-F., AND LITZENBERGER, R.H. *Foundations for Financial Economics*. Amsterdam: North-Holland, 1988.

13. HULL, J. *Options, Futures, and Other Derivative Securities*. Englewood Cliffs, N.J.: Prentice-Hall, 1989.

14. HULL, J., AND WHITE, A. "Pricing Interest-Rate-Derivative Securities," *Review of Financial Studies* 3, no. 4 (1990): 573-92.

15. JACOB, D.P., LORD, G., AND TILLEY, J.A. "A Generalized Framework for Pricing Contingent Cash Flows," *Financial Management* 16, no. 3 (1987): 5-14.

16. MALLIARIS, A.G., AND BROCK, A.W. *Stochastic Methods in Economics and Finance*. Amsterdam: North-Holland, 1981.

17. MILLER, S. "A Continuous Arbitrage-Free Interest Rate Model," *Risks and Rewards* 7 (1990): 4-7. (Society of Actuaries, Schaumburg, Ill.)

18. PEDERSEN, H.W., SHIU, E.S.W., AND THORLACIUS, A.E. "Arbitrage-Free Pricing of Interest-Rate Contingent Claims," *TSA* XLI (1989): 231-65.

19. REITANO, R.R. "Nonparallel Yield Curve Shifts and Durational Leverage," *Journal of Portfolio Management* 16, no. 4 (1990): 62-67.

20. REITANO, R.R. "Multivariate Duration Analysis," *TSA* XLIII (1991): 335-76.

21. RUBINSTEIN, R.Y. *Simulations and the Monte Carlo Method*. New York, N.Y.: John Wiley & Sons, 1981.

22. SHARPE, W.F., AND ALEXANDER, G.J. *Investments*, 4th ed. Englewood Cliffs, N.J.: Prentice-Hall, 1990.

23. VASICEK, O. "An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics* 5 (1977): 177-88.

24. WHEELER, R.E. "Random Variate Generators," *ARCH* 1974.1.

Discussion of Preceding Paper

Michael F. Davlin

I commend Dr. Tilley for writing yet another fine paper for our *Transactions* on the important problem of interest rate risk. This excellent paper successfully achieves his stated goal of presenting the basic concepts of interest rate processes and option-pricing in such a manner that they are easily understood by the many actuaries who have chosen not to pursue the study of modern financial economics and stochastic processes. My own understanding of how a reasonable number of arbitrage-free interest rate paths can be constructed has been much improved by reading this article, and I am anxious to see the author make good his promise (in Section 5.3) of a forthcoming technical note extending the methods he outlines in this paper. I especially appreciate that he points out the room in his model for the important and, in my opinion, unavoidable issues of subjective expectations and subjective risk preferences. A well-written paper always raises questions and observations in the mind of its reader, and this paper is no exception. A few of my thoughts are presented here for the author's consideration.

As the author points out, a key contribution to actuarial practice by financial economics is its emphasis on what the classical economists referred to as the law of one price: in a competitive market equilibrium there cannot exist more than one price for the same good. Financial economists now refer to a generalization of this law as the absence of riskless arbitrage. I am tempted to argue that many actuarial models, those based upon decision-theoretic principles, were (internally) arbitrage-free before our recent encounters with the theories of modern finance in the following sense: identically contingent cash flows within these models were assigned identical prices. The (externally) arbitrage-free condition imposed by the financial economist places additional consistency restrictions or bounds on what a modern actuary should consider to be a reasonable valuation of uncertain cash flows: model prices must match market prices. While I confess to finding this very appealing and intuitive, I also have difficulty adopting it wholeheartedly and without some caution.

Some aspects of my reservations may merely be differences in interpretation, with little or no practical consequences. For example, if I interpret a lognormal distribution of future interest rates as being reasonably representative of my uncertainty about the future course of interest rates, I have little difficulty adopting it for modeling purposes. I may even find myself estimating its parameters from historical prices if I do not have a compelling reason for using some other approach. If, on the other hand, adopting this model means that I have to accept the somewhat metaphysical baggage that often accompanies "scientific" economic modeling—assumed market equilibrium, process stationarity and ergodicity, and a generally positivist outlook on the world—then I become far less comfortable using and accepting these models. In contrast to the author's earlier writings on this subject, this paper admits that there may be validity to reflecting a decision-maker's subjective outlook in modeling. Perhaps I read too much into his remarks. I would be very interested to know how the author interprets his model's interest rate process and whether he believes that the interpretation of a model is important.

Another impediment to my embracing this approach to modeling interest rates is that it does not seem to explain interest rates in a way that gives me any guidance into how current developments might influence the future. In the same sense that Newton's law of gravity does not explain gravity but only describes its effects, the financial economist's models describe the pattern of interest rate movements but do not explain why they move or even why interest rates exist in the first place. As a consequence, these models are not of much use in answering questions that can be meaningful to an actuary. For example, how will future interest rates be affected if recurring proposals to limit the interest deduction on residential mortgages are adopted? How will rates and default premiums differ if deposit guarantee programs are left in place, in comparison to what they will be if these programs are modified? These types of issues evidently do not affect the decisions of a user of the author's model until, ex post facto, they have an observable effect upon interest rates, and only then through a revision in the historically derived estimates of the process parameters or drift terms. Insights on questions such as these frequently seem to be more available in traditional, even non-mathematical or "armchair" economics. This is ultimately not very intellectually satisfying, because I then feel as though my decision-making is based upon two irreconcilable world views.

A more serious reservation concerns the assumption of equilibrium. The author cautions us not to "debate too-strenuously whether small arbitrages can exist . . . because the assumptions are only approximations to reality." I would agree, but what should one do if he does not believe that one or more of the assumptions are not even remotely close to reality?

What reasons do we have to believe that equilibrium is ever attained? Certainly, an unfettered market abhors a free lunch as much as nature abhors a vacuum. But does the fact that market participants have every incentive to discover and consume away any and all free lunches guarantee their success in doing so in the real world? At least since Keynes appeared on the scene, economists have been debating whether equilibrating or disequilibrating forces prevail in a free market. I am not aware that this issue has been resolved either way and am not optimistic that it ever will be. It seems to me that an absence of riskless arbitrage can exist regardless of whether a market is in equilibrium. If this is possible, then it raises the question of how much advantage there truly is to evaluating a strategy or position in the risk-neutral world implied by a constellation of prices in a market that may or may not be in equilibrium, as opposed to judging them with a subjectivist's outlook to the future. Clearly, serious mistakes can result from either approach. I would caution actuaries against concluding that decisions made in the computationally convenient risk-neutral world are by their nature more "objective" or less risky than those made in the real world using methods that reflect preferences and expectations. One of the most attractive aspects of the method presented in this paper is that it seems to allow for both outlooks.

Once again, I am grateful to the author for this stimulating paper and look forward to reading his promised technical note.

Merlin Jetton

Other actuaries, as I do, should appreciate Dr. Tilley's efforts to increase their understanding of interest-rate modeling. He discusses many of the considerations and describes how to create arbitrage-free interest rate scenarios.

My aim is similar, but in a quite different way. I believe we should be aware of the purpose of an arbitrage-free

option-pricing model. Some assumptions that underlie the arbitrage-free hypothesis are discussed. I discuss some theoretical and practical aspects of the interest rate model described by Dr. Tilley.

Such an option-pricing model is designed to calculate *current* market values. The current price (present value) of a financial instrument whose value depends on interest rates obviously must use interest rates to do the present valuing. In other words, such a model is a quite complex present-valuing algorithm. It generates interest rate scenarios, but the scenarios generated are subservient to the primary goal of calculating *current* market values. The scenarios are not intended to answer the question "What might the future be like?" but rather "What is the theoretically fair market value of this or that financial instrument or portfolio given today's yield curve?" Hence, the scenarios that an option-pricing model generates may not be appropriate given a different primary goal, for example, financial projections. (The model user needs to make this judgment.) The scenarios that the option-pricing model does generate must be, in a certain sense, arbitrage-free.

In a free market an entrepreneur who finds an arbitrage opportunity will quickly exploit it and make an extraordinary profit. However, the market's supply and demand conditions will respond to the entrepreneur's actions, and the arbitrage opportunity will soon vanish. In the context of an option-pricing model, this implies that two financial instruments or portfolios that are nearly identical or close substitutes should have identical, or nearly identical, prices. If there is a large enough price discrepancy between the two instruments per the model that it would present an arbitrage opportunity, then the model is not arbitrage-free.

Dr. Tilley discusses interest rates in terms of coupon rates, spot or zero-coupon rates, and forward rates. They are all related. Given a yield curve of coupon rates, one can calculate a corresponding yield curve of spot rates, and then an array of implied forward rates. I wish to make two points about these different kinds of rates:

1. Neither coupon rates nor spot rates make any claim about market interest rates as of some future date. They only portray rates in effect at the moment. On the other hand, forward rates calculated as described can be used to make an inference about market interest rates as of future dates. If you assume that the expected spot rate at some future date depends on a

forward rate that exists today, then you are making an assumption.

2. All these different kinds of rates can be used, in somewhat different ways, to calculate the present value of a cash flow. For ease of discussion, I refer to the three different ways as the yield-to-maturity method (YTMM), the spot rate method (SRM), and the forward rate method (FRM). The last method is quite different from the other two. Using the FRM starts with a spot rate as of now. This permits calculating the present value of a cash flow that will (or might) occur at the end of the period coterminous with the initial shortest-period spot rate. To calculate a present value of a cash flow that will (or might) occur beyond the end of such period, one links together inferred forward rates as of one or more future dates. An assumption is made about future interest rates.

I have read a limited amount of the literature on arbitrage-free interest rate models. But it seems to me, and it is quite clear in Dr. Tilley's model, that such a model has these features:

1. The starting yield curve is taken as given, and one would presumably use the current real-world yield curve.

2. The FRM is used to calculate all present values. Each linked forward rate covers an equal amount of time, and the forward rates are as of the future dates at which the model assumes interest rates can change.

3. For the model to be arbitrage-free, the following conditions must hold. Using the FRM, calculate the present value of $1 payable at any future date n years from now along each path of interest rates the model projects, and then average such present values. Then calculate the present value of the same $1 using the n-year spot rate based on the starting yield curve. These two present values must be equal. It is presumably this forced condition that calls for all the intimidating mathematics to which Dr. Tilley refers.

Because of this third feature, the starting yield curve determines the average trend of future interest rates, at least of the shortest term interest rate, that the model will stochastically generate. If the starting yield curve is upward-sloping, the model will generate on average higher interest rates as it proceeds forward in time. If the starting yield curve is downward-sloping or flat, the model will generate rates that on average trend downward or flat as it proceeds forward in time. If the

starting yield curve is humped, the model will generate on average higher interest rates, and then lower interest rates, as it proceeds forward in time. In other words, the model will "walk up" or "walk down" the starting (forward rate) yield curve. Of course, this affects the likelihood of options being exercised and the values the model places on such options.

The arbitrage-free model almost exclusively uses the FRM as the present-valuing method. Is it the best method for present-valuing all cash flows? Is it better than the SRM for calculating the present value of a Treasury note or bond? Is it better than the YTMM for calculating the present value of a Treasury coupon-paying note or bond? Let us consider an example. Suppose we want to value at time 0 a five-year zero-coupon Treasury note, with a face value of $100, whose current market price is $68.05. According to the arbitrage-free hypothesis, we must be able to calculate the present value, using the FRM, of that $100 along each of the future interest rate paths, average the results, and get $68.05. It strikes me as odd that the FRM calculates different present values for the different paths when the present value of the $100 using the SRM for any path is $68.05. I guess it doesn't matter if the average is the only concern. However, the FRM would give an invalid measure of the variance of returns; the SRM would not.

According to the arbitrage-free hypothesis, the expected return (the mean of an assumed probability distribution) for a given holding period should be equal for bonds of different maturities. For example, the expected total return for investing in a two-year U.S. Treasury for one year should equal the expected return of investing the same amount in a one-year Treasury, even when the yield-to-maturity for the former is greater. Suppose you were to use the model, but you do not force it to "walk up the forward yield curve"; rather, you set the drift parameter so that interest rates would on average not drift upward. That would be a violation of the arbitrage-free hypothesis. The two-year Treasury, vis-a-vis the one-year Treasury, presents an arbitrage opportunity. Its present value, averaged over all paths, exceeds the present value, averaged over all paths, of the one-year Treasury. Is this truly a riskless arbitrage opportunity or is it a quirk of using the FRM? Consider the following:

1. As stated earlier, a necessary condition for a riskless arbitrage opportunity is that it can be exploited nearly instantaneously. The arbitrage opportunity that you want a model to avoid is inconsistent prices

based on current interest rates, not current interest rates and inferred future interest rates as well.

2. Amazingly, this apparent riskless arbitrage opportunity would vanish if the SRM were used. It does not depend on inferred interest rates.

It was just suggested that the model will not replicate current market prices when not set in the arbitrage-free mode. This is generally true. Subsequent changes in market price implied by using a non-arbitrage-free drift parameter are captured in the present value at time 0. This is again the result of relying on the FRM in contrast to the SRM. The SRM would not do this.

If one is going to use this model regularly, there are further implications to be aware of. Suppose the following. One day when the yield curve is quite upward-sloping, you use the model for a particular set of assets and liabilities. Then a few weeks later, you use the model again for essentially the same assets and liabilities. Not much has changed in the few weeks, except that the short end of the yield curve has shifted upward, making it less upward-sloping than before. The results from the second time will differ markedly from those the first time. This is not a concern if you are interested only in current market values. However, if you are concerned with other results, you will probably question the model and the assumptions you made. You recognize that some of the results produced using the model are fairly crude estimates anyway. You know very well that yield curves shift often, and in two weeks the curve may be back where it was when you ran the model the first time. Would you be satisfied with your results?

I do not think that the starting yield curve has to dictate the average trend of future interest rates. Otherwise it is implicitly assumed that the forward rates derived from the starting yield curve do not contain any risk premiums. I find that assumption very unrealistic. How is it that spot rates contain risk premiums, but forward rates inferred from them do not? It would be more realistic to assume that these forward rates do contain risk premiums and that these risk premiums generally decline and eventually vanish as each forward rate comes ever closer, with the passage of time, to becoming a spot rate. This contrary assumption, of course, introduces a judgmental element. The amounts of these risk premiums are not easily observable in the marketplace, so different people will make different judgments about their size and thereby determine different average trends of future interest rates given the same starting yield curve. On the other hand, I do not believe we

should equate attaining objectivity of current prices—by assuming zero risk premiums—with attaining objectivity of future prices (interest rates).

The instruments the arbitrage-free model was developed to price are "derivatives" of other instruments that are close substitutes. They are traded in active and efficient markets. The possibility of arbitrage opportunities being exploited, if they do arise, is more likely. The market value of such options is quite sensitive to shifts in the yield curve in the short run. The time to expiration of most such options is rather short—two years or less. On the other hand, the financial instruments an actuary deals with often contain options that do not have close substitutes or are not as efficiently priced. There is no active and efficient market. Also, such financial instruments often have much longer lives—10 years or more. This suggests that maybe the model needs revision when applied to these somewhat different instruments.

In an earlier paper[9], I said that bias, in some form or another, seems to be an unavoidable feature of stochastic interest rate generation models. I still believe that and that the arbitrage-free model is no exception. As Mr. Miller said in his discussion of that paper, "... our goal should be to know the implications of any bias and to avoid making bad decisions because of it." He also said that the "bias" in option-pricing models comes from assuming that a model that works well for one purpose will work well for all purposes.[10] Another bias in the option-pricing model is its devotion to the FRM for calculating present values when other proven methods may be better.

My final point is that I do not believe jointly using two interest rate models, one of which is arbitrage-free and the other not, would necessarily be incompatible. I will attempt to convey this notion with an imaginary conversation. It is between two actuaries, Gus and Norb. Gus is uncomfortable using an arbitrage-free model for a particular task at hand. Norb (his coworkers sometimes pronounce his name like "no arb") works with Gus and is enthralled with his arbitrage-free model. (He often says, "It's really elegant, a brilliant application of the pure expectations hypothesis.")

Gus: I don't believe using your option-pricing model is the right tool to use here.

Norb: But if we don't use the arbitrage-free model, then we cannot accurately value the embedded options in the assets or the liabilities.

Gus: Please, Norb, bear with me. Look, I have a solution to our dilemma. You walk along my interest rate paths with me. At every stop you can go off and do your thing. You take the yield curve we have at that time, make whatever assumptions you want to make about future interest rates and bring me back the appropriate values for the embedded options. Then you'll again walk along my path with me to the next stop, where you'll repeat the process and I will again wait. It doesn't bother me that the assumptions you make to value the options don't match the assumptions I make about the paths we will walk together.

Satisfied, Norb walks off to his supercomputer to gear it up for the task.

Hal W. Pedersen[11]

Dr. Tilley has written an interesting paper that clearly explains the fundamental ideas of term structure modeling. Furthermore, the paper is written carefully with due regard for the difficulties that an "actuarial layman" might face in an endeavor to understand and implement a "stochastic interest rate generator." I admire the paper for its successful exposition, with the use of but one formula, of such a technical subject. The objectives of the paper, as stated, make it inevitable that it must leave some issues by the wayside. I have some questions, the answers to which will assist me in better understanding not only the ideas of this paper but also the gap between theory and practice. At the same time, I hope that my remarks will provide an alternative point of view on some techniques of this paper.

The discussion begins with a description of the martingale approach to the modeling of the term structure of interest rates. This is followed by a specific example based on Brownian motion. Some questions concerning the model of this paper are then posed.

Martingale Approach to Term Structure Modeling

One standard approach to the modeling of the term structure of interest rates is known as the martingale approach. The work of Artzner and Delbaen [1] is a thorough, though demanding, reference on this topic. The book of Duffie [3, chapter 7] can also be consulted for this material. A different approach has recently been developed by Heath, Jarrow, and Morton [5]. As I agree with Dr. Tilley that continuous state space models are

preferable, the approach to modeling that I describe here is set up as such. Essentially, a model of the term structure may be prescribed by Equation (1).

Basic to the analysis is a probability space $(\Omega, \mathfrak{F}, P)$ equipped with a filtration $\{\mathfrak{F}_t\}$. The filtration is an increasing family of sub σ-algebras of $\mathfrak{F}$; by increasing is meant $\mathfrak{F}_s \subset \mathfrak{F}_t$ for $s < t$ and the filtration is interpreted as the way in which the information in the market evolves or is revealed through time. The filtration is often referred to as the information structure. Also given is a spot rate process r, which is analogous to the one-period rate in Dr. Tilley's analysis. The spot rate process should be everywhere positive to avoid undesirable interest rate behavior such as negative interest rates. For $0 \le t \le T$, let $P(t,T)$ denote the price at time t of a default-free zero-coupon bond that pays 1 at time T. Note that $P(T,T) = 1$. We now fix a risk-neutral probability measure, which is denoted by Q. Note that we are choosing this measure and that consequently the procedure that we are following corresponds to approach (ii) that Dr. Tilley describes on page [294]. A model of the term structure of interest rates can now be prescribed by the formula

$$P(t, T) = E^Q\left[\exp\left(-\int_t^T r_u du\right)\middle|\mathfrak{F}_t\right], \quad (1)$$

where E_Q denotes expectation with respect to the probability measure Q. If it is desired to reduce this formula to discrete time, then this translation can be achieved by replacing the integral with the appropriate summation. It follows from Equation (1), by the law of iterated expectations, that for $0 \le s \le T - t$ the bonds in the model will satisfy

$$P(t, T) = E^Q\left[\exp\left(-\int_t^{t+s} r_u du\right)P(t+s, T)\middle|\mathfrak{F}_t\right], \quad (2)$$

which can be interpreted as a consistency condition. In practical applications, the model prescribed by (1) will not be arbitrage-free. The problem is that the initial bond price function that results from this equation, namely, $P(0,T)$, usually will not agree with the bond prices currently in the market. One approach that can remedy this problem has been documented by Dybvig [4]. It should also be noted that Dybvig shows how to translate a closed formula for the value of an option in the original model to a closed formula for the value of the option in the adjusted model.

The idea is to add a deterministic function, say h, to the spot rate process, which adjusts the model to fit the observed bond prices. As Dybvig [4, p. 6] notes, this perturbation technique is such that, " [t]he variance assumption [of the perturbed model] is as reasonable as it is in the chosen model." Suppose that the observed bond prices are given by the function $f(T)$. Then the problem of fitting the observed bond prices is equivalent to choosing the function h such that for every T

$$f(T) = E^Q\left[\exp\left(-\int_0^T (r_u + h_u)du\right)\right]. \quad (3)$$

Since the function h is deterministic, Equation (3) can be expressed as

$$f(T) = \exp\left(-\int_0^T (h_u du)\right)E^Q\left[\exp\left(-\int_0^T (r_u du)\right)\right]$$

$$= \exp\left(-\int_0^T (h_u du)\right)P(0, T), \quad (4)$$

where $P(0,T)$ is the bond price function at time 0 that the original model generated and $f(T)$ is the observed bond price function that we want the model to generate. Now it is relatively straightforward to solve for the function h. We can then employ the model given by Equation (1) with the spot rate process $r + h$ in place of r to obtain an arbitrage-free model of the term structure.

Let us now describe the option-pricing theory in a model such as (1). The term "contingent claim" is a general term that encompasses the term "option." A contingent claim payable at time t is any non-negative random variable that is measurable with respect to $\mathfrak{F}_t$. For example, a European option with maturity t and strike price K that is written on a bond with maturity T, where $T > t$, is a contingent claim payable at time t with contingent payoff equal to the maximum of 0 and $P(t, T) - K$. Subject to the same assumptions that Dr. Tilley makes, namely, no arbitrage and completeness, the price at time 0 of a contingent claim payable at time t, say, X, is characterized by the formula

$$\text{price of X at time } 0 = E^Q\left[\exp\left(-\int_0^t (r_u du)\right)X\right], \quad (5)$$

where the risk-neutral measure Q is the same Q that appears in the formula for the bond prices given in Equation (1). More generally, the price of the same claim at any time, $s, s < t$, is given by the equation

$$\text{price of X at time } s = E^Q\left[\exp\left(-\int_0^t r_u du\right)X\middle|\mathfrak{F}_s\right]. \quad (6)$$

Taken together, Equations (1) through (6) represent one standard approach to the construction of arbitrage-free term structure models and the pricing of interest-rate-contingent claims within these models.

There is a fundamental difference between term structure modeling and the option-pricing within the model despite the fact that the expectation formulas in Equations (1) and (6) are similar. Equation (1) is the definition of the bond price processes. In other words, Equation (1) prescribes the bonds for the model. In contrast, Equation (6) evaluates the price of the contingent claim. An important assumption behind Equation (6) is that the cash flows from the contingent claim can be replicated by an appropriate self-financing portfolio of the assets, and this allows us to price the contingent claim by our knowledge of the prices of the assets in the replicating portfolio. In the succinct terminology of Dr. Tilley, Equation (6) represents the "relative price" of the contingent claim.

The pricing of contingent claims in the approach to term structure modeling that I have described here is really a two-step procedure. The first step is a prescription of the bond prices, which is made by fixing a risk-neutral probability measure and a spot rate process and then defining the bond prices through Equation (1). The second step is to employ these bond price processes and then to price the contingent claims through Equation (6). The recent paper [5] objects to this two-step procedure. There are sound theoretical reasons for their objection, and consequently there are good reasons to perhaps favor the approach that is taken in [5], as Dr. Tilley has alluded to at various points of his paper. Nevertheless, this two-step procedure is a time-honored approach in the finance literature and provides a benchmark for understanding other models such as [5].

Last, Equation (5) underscores the important point on page [307] about the role of the one-period rates. For a concrete example to illustrate this point, consider a European option with maturity t and strike price K that is written on a bond with maturity T, where $T>t$. This contingent claim provides a cash flow equal to Max[0, $P(t, T)-K$], which is clearly determined by the bond price and not the spot rate, and by Equation (5) we also see the role of the spot rate in the valuation of this claim, because

price of option =

$$E^Q\left[\exp\left(-\int_0^t (r_u du)\right)\text{Max}[0, P(t, T) - K]\right].$$

An Example Based on Brownian Motion

Let us now mention a concrete case of the model [Equation (1)] that is known to be arbitrage-free and complete and thus meets Dr. Tilley's basic assumptions in Section 4.1. This specialization also will provide an example of what could come out of a path simulation of the model (1). To this end, let us now specialize the information structure to be a Brownian filtration and the spot rate process, r, to be defined by a stochastic differential equation

$$dr_t = \mu_t dt + \sigma_t dW_t.$$ (7)

In Equation (7), W is to be understood as a vector of independent Brownian motion processes, $W = (W^{(1)}, W^{(2)}, ..., W^{(N)})$, and σ is understood to be an N-dimensional vector process. In the case of $N = 1$, we have a one-factor model and in the case of $N > 1$, we have a multifactor model. For a more general formulation of the multifactor case, one could consult Duffie [2, p. 137]. In this setting, Artzner and Delbaen [1, section 2.2] show that the bond price processes that result from Equation (1) are of the form

$$\frac{dP(t, T)}{P(t, T)} = \mu_t^T dt + \sigma_t^T dW_t.$$ (8)

Please note that the differential in this equation is to be understood in t, for $t < T$.

If we choose a number of bonds, indexed by maturity times, say, $T_1, T_2, ..., T_K$, then using Equation (8) we can write

$$\begin{bmatrix} \dfrac{dP(t, T_1)}{P(t, T_1)} \\ \cdot \\ \cdot \\ \cdot \\ \dfrac{dP(t, T_k)}{P(t, T_k)} \end{bmatrix} = \begin{bmatrix} \mu_t^{T_1} \\ \cdot \\ \cdot \\ \cdot \\ \mu_t^{T_k} \end{bmatrix} dt$$

$$+ \begin{bmatrix} \sigma_t^{1T_1} & \sigma_t^{2T_1} & ... & \sigma_t^{NT_1} \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \sigma_t^{1T_k} & \sigma_t^{2T_k} & ... & \sigma_t^{NT_k} \end{bmatrix} \begin{bmatrix} dW_t^{(1)} \\ \cdot \\ \cdot \\ \cdot \\ dW_t^{(N)} \end{bmatrix}$$ (9)

which gives a matrix representation for the interrelationships between the rates of return on the various bonds. Equation (9) also exhibits the correlation structure between these various bonds.

For this model, the technique in Equation (3) is precisely the addition of a time-dependent drift to the spot rate process. This can be seen immediately from the stochastic differential equation for the spot rate r. Thus, it appears that there may be some analogy between the adjustment procedure that Dr. Tilley employs and the type of adjustment procedure in Equation (3).

Dr. Tilley states on page [290] that, "a *single* stochastic factor—the short-term rate of interest—drives the dynamic of the *entire* yield curve, resulting in perfect correlation of yield movements across the curve." Some clarification of this point in the context of the general model that I have described may be appropriate. A model that is prescribed by Equation (1) is such that all the interest rate dynamics are driven *through* the spot rate process r. However, in the case in which there is more than one independent Brownian motion that is driving the spot rate, then we do not have perfect correlation of the yield curve movements. This fact is emphasized by Equation (9).

In the case where r is driven by only one Brownian motion, I wonder if the following interpretation is what Dr. Tilley had in mind? It is convenient to drop the superscript and denote the Brownian motion by W. Equation (8) shows that the returns on all bonds are locally perfectly correlated because dP/P are linear functions of one another across all bonds. An application of Itô's lemma to (8) shows that for $s < t$

$$d[\log P(t, T)] = \left[\mu_t^T - \frac{1}{2}(\sigma_t^T)^2\right]dt + \sigma_t^T \, dW_t.$$

A discretized version of this equation would take the form

$$\log\left[\frac{P(t+1, T)}{P(t, T)}\right] =$$

$$\left[\mu_t^T - \frac{1}{2}(\sigma_t^T)^2\right] + \sigma_t^T (W_{t+1} - W_t). \qquad (10)$$

This equation could be recast in the notation of Section 5.2, which is

$$r_{k,t} = \frac{-1}{k-t}\log P(t, k).$$

In particular, Equation (10) shows that there is perfect correlation of yield movements across the curve. Indeed,

the yield movements are perfectly correlated in the usual statistical sense because these movements are all linearly related to the same random variable, $W_{t+1}-W_t$. In a fashion similar to Equation (10), a discretized version of the multifactor case could be written down, but the yield movements now will be imperfectly correlated. For a Brownian motion process, the random vector $W_{t+1}-W_t$ has a normal distribution. However, the actual distributional properties of a discretized version of this multifactor case will depend on the properties of the coefficients as well.
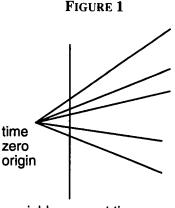
Some Questions

Can the model for generating interest rates described in this paper be viewed as a path-by-path simulation (or sampling) of a model that is prescribed by a formula such as Equation (1) and that is then adjusted to fit the observed bond price data? If this is the case, then one could feel at ease with this model because this model could be regarded as being embedded in a standard arbitrage-free framework. Also, what exactly is the theory behind this model? Furthermore, if this simulation interpretation is correct, is it possible to make rigorous the convergence argument, based on equal probability samples, that appears on page [294]? Also, I wonder whether it is possible to make rigorous the part of this assertion that says, "[t]hus, when all possible paths are represented, solving for the risk-neutral drift of the stochastic process is equivalent to establishing the risk-neutral probabilities." If such an interpretation is not correct, then can Dr. Tilley offer a precise alternative interpretation that would also explain what is meant by the remark on page [290], "... by way of calculations performed on paths sampled from a continuous model?" What is the relationship of the model that is prescribed in this paper to the class of models that are given by Equation (1)? There should be some relationship here because Dr. Tilley states on page [292] that "I now show how this description of the risk-neutral world can be used to construct an arbitrage-free model of interest rates."

A simulation of the type of model described by Equation (1) under an adjustment procedure such as Equation (3) seems to be analogous to what is done in Section 4.2 and to what is suggested on page 102[*sic*]. The multifactor discretization will not be analogous to the most general model in Section 5.2, because the general model of this section prescribes all the rates simultaneously and because the drift adjustment factor could be stochastic. However, the multifactor discretization

does seem to be analogous to Dr. Tilley's model in the case in which it is based only on the "key rates (page [296])" and the drift adjustment factor is deterministic. In brief, if we were to discretize this model and then simulate the discretized dynamics, it appears that we would be doing something like what Dr. Tilley is doing. Although this analogy with Dr. Tilley's generator is not perfect, it is suggestive.

In some sense can we regard Dr. Tilley's model as a simulation, perhaps with a discretization, that is embedded within a master model of the term structure? If so, perhaps Dr. Tilley's model is best interpreted as a way of generating a consistent interest rate scenario, as a sampling from a full term structure model, and then applying the pricing techniques that the full model tells us are correct. Under such an interpretation, this approach would not be taken as a literal model of the term structure but rather as a way of using a master term structure model to price complicated cash flows in a computationally feasible way.

Although I do not fully understand the ideas behind Dr. Tilley's model, the following description of the procedure seems to be correct, and it summarizes a major point of confusion that I have. Indeed, once the interest rates are generated by the model, it appears that one might view the situation as follows. At time 0 one does not know which path one will be on, and so in some sense the yield rates are stochastic at this time. However, as soon as you move ahead to time 1, the yield rates are then deterministic because you will be on the same path of yield curves for the rest of the time. Figure 1 illustrates the point. Loosely speaking, once you take the first step to the state of the world at time 1, then you are in a deterministic world. This would then seem to lead to arbitrage because, for instance, the forward rates imply definite restrictions if the next yield curves are known with certainty.

"[T]hus, constraining the process to be arbitrage-free is tantamount to setting its time dependent drift so that all zero-coupon bonds are priced at values equal to those derived from the exogenously specific initial yield curve" (page [293]). Is this really an appropriate notion of an arbitrage-free model? Does the more fundamental notion of such a model being arbitrage-free involve being free of arbitrages that arise by setting up certain riskless hedges through the trading of bonds? Since there is no trading of bonds possible in this path-based model, it seems unclear how the usual notion of arbitrage-free would be understood.

FIGURE 1

yield curves at time one

Another question is how one is to value contingent claims (options) that have early-exercise provisions. The following two quotations from page [290] seem to suggest that Dr. Tilley is not advocating the usual arbitrage-based methods for valuing contingent claims with early-exercise provisions, "[w]ithout a lattice (whether connected or not), backward induction is not possible. From a purely mathematical viewpoint, it is difficult to construct *optimal* exercise strategies for many option problems by doing calculations on interest rate paths sampled from a continuous model," and "[t]he behavior of these people, as to their strategies for rational (if not mathematically optimal) exercise of the options they hold, can be modeled sufficiently accurately that the options are valued properly by way of calculations performed on paths sampled from a continuous model."

My first point of confusion is what is meant by the word "difficult." Backward induction is the only method that I know of for constructing the optimal exercise strategies that are necessary to evaluate contingent claims that have early-exercise provisions. Because the model in this paper does not permit backward induction, I do not know what technique Dr. Tilley has in mind.

I also point out that it is not possible to project the optimal cash flows from such a contingent claim on a path-by-path basis. All possible future events must be taken into account to make this optimal projection. In other words, at each epoch all the future possibilities that may develop from the current state of the world must be known, but such is not the case for this model because there is no evolution structure available. For example, even though any particular contingent claim that has an early-exercise provision may be in the money, it does not mean that the claim holder will want to exercise it, although he may wish to do so in some

situations. (A good example of this dual phenomenon is the American put option.) The exercise decision will depend on the future states of the world (hence future values of the claim) in relation to the current state of the world. Hence, it seems impossible to value certain contingent claims by this model in accordance with the usual arbitrage-based techniques. In fact, it seems this model can only properly price claims whose cash flows at each time along a path are determined by the yield curves along that path only. (By properly, I mean a price that is consistent with the benchmark of arbitrage-based pricing.) A particular subclass of such claims would be those claims whose cash flows at each time are determined by the past history and the current yield curve only. For instance, the model in the paper would price interest rate caps and floors nicely. Is it possible to quantify the bias in pricing contingent claims that have early-exercise provisions that results from using the approach that Dr. Tilley is suggesting as compared to the arbitrage-based optimal exercise procedure?

I am also unclear on how to reconcile this second quotation from page [290] with the following quotation, which is taken from [2] in the context of a particular example but which is representative of a general principle in pricing by arbitrage-based methods, "[i]t still might seem that we are depending on rational behavior by the person who bought the call we sold. If instead he behaves foolishly and exercises at the wrong time, could he make things worse for us as well as for himself? Fortunately, the answer is no. Mistakes on his part can only mean greater profits for us."

Note that for the class of models that I have described, the optimal exercise strategies can be computed because we can project the optimal cash flows since there is an information structure (filtration) present at each state of the world. Thus, in theory, contingent claims that have early-exercise provisions can be priced by arbitrage-based methods. In practice, these arbitrage-based prices have to be computed numerically, and this is a most troublesome task.

I thank Dr. Tilley for an insightful paper that has furthered my understanding of some of the difficult issues involved in the practical application of term structure theory. I also appreciate the lesson in the art of communicating difficult concepts without relying on formulas. However, it would be most interesting to see some pictures of the yield curves, with the accompanying parameters, that are generated by the model of Section 5.2. This would be helpful in appreciating the wide range of yield curve dynamics that the model is capable of producing.

References

1. ARTZNER, P., AND DELBAEN, F. "Term Structure of Interest Rates: The Martingale Approach," *Advances in Applied Mathematics* 10 (1989): 95-129.
2. COX, J., ROSS, S., AND RUBINSTEIN, M. "Option Pricing: A Simplified Approach," *Journal of Financial Economics* 7 (1979): 229-63.
3. DUFFIE, J.D. *Dynamic Asset Pricing Theory.* Princeton: Princeton University Press, 1992.
4. DYBVIG, P. "Bond and Bond Option Pricing Based on the Current Term Structure," Working paper, Washington University, Olin School of Business, 1988.
5. HEATH, D., JARROW, R., AND MORTON, A. "Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Valuation," *Econometrica* 60 (1992): 77-105.

Author's Review of Discussions

James A. Tilley

I thank Mr. Davlin, Mr. Jetton, and Mr. Pedersen for their discussions of my paper. I find it most convenient to respond to their questions by author.

Mr. Davlin's Discussion

I believe that an interest rate model should capture how the world actually works. One should do the best job possible to discover the real world's interest rate dynamics. If one observes fat-tailed distributions, tendencies to mean reversion, stochastic rather than deterministic volatility, liquidity premiums, and so on, one should build them into the model. Subjective views of the trend of interest rates and the shape of the yield curve also should be incorporated into the model. Of course, any model will end up being only an approximation to reality, but it should be the best one can build. The most important use of such a model is to evaluate risk/reward tradeoffs characteristic of all portfolio decisions.

Once such a model is available, it also can be used to calculate consistent relative prices among various instruments whose cash flows either are fixed or depend only on interest rates. For this purpose, and probably

for this purpose only, the drifts of the interest rate processes must be adjusted so that the desired relative prices can be derived properly by way of straightforward expected-value computations. Once one has determined the fair market prices for all instruments, one should return to the real-world model (the one with *unadjusted* drifts) to evaluate the relative attractiveness of various portfolio strategies.

I share Mr. Davlin's yearning for models that "explain" sources of interest rate and spread changes. Such models are the subject of econometrics, a discipline in which mathematical economists attempt to codify the fuzzy ideas of armchair economists. In general, the results of such efforts have been rather unsatisfying. Even a satisfactory econometrics model for interest rates would be probabilistic, not deterministic. The purpose of my paper was not to take on the challenge of building a better econometrics model. But I would like to say a little more about the question of the "causes" of the stochastic dynamics of interest rates.

Consider the case of the Brownian motion of a particle in a liquid. One knows that the particle moves within the liquid because it is buffeted by molecules of the liquid—in other words, one understands the "causes" of the particle's motion. One even understands the physics of the interaction among molecules of the liquid and their resulting motion. Nevertheless, the best description, indeed, the only useful description, of the dynamics of the particle is a random movement governed by a particular stochastic differential equation. In the case of the particle, one can "derive" the stochastic equation of motion, but in the case of interest rates, one has to "observe" behavior in the real world and then select a satisfactory equation of motion empirically. I believe it is the "empirical" approach to which Mr. Davlin objects on philosophical grounds. Yet, actuaries often utilize empirical approaches in their traditional areas of work—they observe the age dependency of the forces of mortality and morbidity, and they observe the probability distributions of losses that arise from the violence of the earth, air, water, and fire.

Mr. Davlin poses good questions about market equilibrium, questions that many good minds will ponder for years to come. Related to his comments in this area, let me stress once again that I believe that *decisions should be made in setting of the real world* using methods that reflect subjective preferences and expectations—I do not believe that there is necessarily any advantage to evaluating a strategy in the setting of the risk-neutral world.

However, I do know that *fair market prices can be calculated easily in the setting of risk-neutral world.*

Mr. Jetton's Discussion

My comments on Mr. Davlin's discussion indicate that I do not agree with what I believe to be Mr. Jetton's view that "[interest rate] scenarios are not intended to answer the question 'What might the future be like?'."

I also disagree with the thrust of Mr. Jetton's statement that "neither coupon rates nor spot rates make any claim about market rates as of some future date ... [but] forward rates ... can be used to make an inference about market interest rates as of future dates." This ostensibly innocuous statement is the opening salvo of his view that the use of forward rates to perform discounting calculations involves assumptions (implicit or otherwise) that are not needed by or are not shared by other discounting methods. I find many of Mr. Jetton's statements about calculating present values using current-coupon yields (his label is YTMM), spot rates (his label is SRM), and forward rates (his label is FRM) to be confusing, misleading, or inaccurate. I offer a detailed critique of a few of Mr. Jetton's examples below.

As indicated at the end of Section 2 of my paper, the sets of current-coupon yields, spot rates, and forward rates contain *equivalent information*. Any one set of yields or rates is sufficient to derive the other two sets. Depending on the situation, there may be a natural set to use, but all carry identical information. These are mathematical statements of fact, not opinions.

When cash flows are non-interest-sensitive, it makes absolutely no difference whether one uses yields to maturity, spot rates, or forward rates to calculate present values, provided that the calculations are performed properly and accurately. For this situation, generating interest rate scenarios is unnecessary. Computing the present value of a non-interest-sensitive cash-flow stream by using the relevant spot rates is not only correct, but also by far the simplest method. However, when cash flows are interest-sensitive, the only practical choice for discounting the cash-flow stream to a proper present value is the "expected present value" method, invented by others and described by me in the paper. For this situation, the SRM as described by Mr. Jetton will give the *wrong* answer. Whether one constructs the random generator for interest rates based on current-coupon yields, spot rates, or forward rates is partly a matter of taste. Thus, it is not proper to single

out forward rates and claim that interest-rate-contingent cash flows must be discounted by something called an FRM. Instead, it is proper to state that interest-rate-contingent cash flows must be discounted by the expected-present-value method, and that the paths of future interest rates can be sampled from stochastic processes relating to current-coupon yields or spot rates or forward rates. Specifying the stochastic processes in terms of forward rates makes it easy to constrain the forward rates from ever becoming negative, a statement that does not apply when the stochastic processes are specified in terms of either spot rates or current-coupon rates.

After reading Mr. Jetton's comments, I think that it is important to stress a point about the expected-present-value method that was made in my paper and that was reinforced by a comment made by Mr. Pedersen in his discussion. *The one-period rates of interest play a special role in the expected-present-value algorithm.* Only the one-period rates are used to discount cash flows along a path. Yields of bonds with maturities greater than one period affect the arbitrage-free price of a stream of interest-rate-contingent cash flows only to the extent that the amounts of the cash flows depend on those yields. Note that a one-period rate is a current-coupon yield, a spot rate, and a forward rate. Certainly, it seems artificial to claim that the expected-present-value algorithm is a "forward rate method" (FRM), as Mr. Jetton states.

In giving an example of pricing a five-year $100 face amount zero-coupon bond, Mr. Jetton states that the present value is $68.05 for any path if one uses the SRM, whereas using the FRM leads to different present values for each path, these different values only *averaging* to $68.05 over the full set of arbitrage-free paths. He then states "I guess it does not matter if the average is the only concern. However, the FRM would give an invalid measure of the variance of returns; the SRM would not." For Mr. Jetton's example, the zero-coupon bond has a single fixed cash flow, not interest-sensitive cash flows. Thus, as I stated above, using the five-year spot rate to discount the bond's single cash flow is the most straightforward way to obtain its price. Also, as I stated above, one does not need to generate an arbitrage-free set of interest rate paths to price this bond, but if one chooses to generate such a set of paths, the expected-present-value method will lead to the correct price.

Where Mr. Jetton errs is in comparing the returns, and hence the variances of the returns, for two *different* investment strategies. The first strategy is to buy the five-year zero coupon bond at time 0 and to hold it for five years until its maturity—the terminal wealth is $100, regardless of the actual path of interest rates followed. The second strategy is to invest $68.05 at time 0 in a one-period bond (for example, a three-month Treasury bill if the period between epochs is one quarter of a year), and to keep rolling that investment at its maturity into successive one-period bonds until the five-year horizon is reached. For this second strategy, the terminal wealth depends on the path that interest rates follow. For some paths, the terminal wealth is greater than $100, while for others is it less than $100. In fact, the terminal wealth *generally does not* average $100, and *definitely does not* average $100 if the full set of paths is arbitrage-free, a very surprising, but correct, result. The variances of the returns are different for the two strategies—zero for the first strategy and positive for the second strategy—because the strategies are different!

Mr. Jetton's statement that "the expected return for a given holding period should be equal for bonds of different maturities" is not generally true. The expectations that must be equal relate to *prices* (that is, present values), not to returns (that is, future values). Only *locally,* that is, over a *one-period horizon,* must expected returns for bonds of different maturities be equal. This is a very commonly misunderstood point that is almost always discussed too casually in the literature on interest rate processes.

Mr. Jetton offers a second example to highlight the advantages of what he has called the SRM and the difficulties encountered by what he has called the FRM. The example involves the returns from two different investment strategies: (1) buy a one-year Treasury and hold it for one year, and (2) buy a two-year Treasury and hold it for one year. Mr. Jetton properly describes a real-world situation in which the drifts of the interest rate processes are such that the second strategy has a greater expected return than the first strategy. He then states: "The two-year Treasury, vis-a-vis the one-year Treasury, presents an arbitrage opportunity. ... Is this truly a riskless arbitrage opportunity or is it a quirk of using the FRM? ... Amazingly, this apparent riskless arbitrage opportunity would vanish if the SRM were used. It does not depend on inferred future interest rates."

The situation that Mr. Jetton poses is realistic. Most investors consider it to exist every time the yield curve

is sufficiently positively sloped. However, his analysis is flawed in several respects. First, the situation has nothing whatsoever to do with SRMs or FRMs—how Mr. Jetton makes this connection I do not understand. Second, there is nothing wrong with a situation in which the two strategies have different expected returns. Third, *the opportunity available to the investor in this example is not a riskless arbitrage*. By shorting the one-year bond and using the proceeds to purchase the two-year bond, one will lose money if interest rates rise a lot by the end of the year—there is nothing riskless about the arbitrage!

In my opinion, other misleading, if not incorrect, statements made by Mr. Jetton in connection with present-value calculations are: "This is again the result of relying on the FRM in contrast to the SRM. The SRM would not do this" (page [305]) and "Another bias in the option-pricing model is its devotion to the FRM for calculating present values when other proven methods may be better" (page [306]).

I agree with Mr. Jetton in that "I don't believe it has to be the case that the starting yield curve dictates the average trend of future interest rates." Portfolio and business decisions should be based on a real-world calibration, of the interest rate model, in which the drift of interest rates is likely to be somewhat, if not fully, independent of the starting yield curve. However, for purposes of calculating the prices of interest-rate-contingent cash-flow streams in a straightforward fashion that produces answers consistent with the observed prices of Treasury bonds, one must adjust the drift of interest rates based on the starting yield curve.

Mr. Jetton's character Gus essentially has it right: Norb is allowed to "do his thing" by using an arbitrage-free model to value embedded options. Gus then uses this information at each point along his own real-world paths of interest rates to make whatever business decisions he thinks are appropriate. By the way, Norb does not actually need a supercomputer—a powerful PC or a good workstation will suffice.

Mr. Pedersen's Discussion

The adjustment procedure used by Mr. Pedersen in his example is the type that I have described in my paper. His clarification of what is meant by a multifactor process is very helpful. The perfect local correlation of yield curve movements across the yield curve in the

case of a single Brownian motion is exactly what I had in mind.

The interest rate generator that is described in my paper can be viewed as a path-by-path simulation of a model described by Mr. Pedersen's Equation (1) and that is then adjusted to fit the observed bond price data. Apart from the discretization of time, I prefer to think of the models I have described as being exactly those described by Mr. Pedersen. *The selection of a finite sample of paths from the model is merely a computational device to price contingent claims*. Different samples of P paths each will lead to different prices for a given contingent claim. The probability distribution of these sample estimates has a variance that tends to zero as P tends to infinity. The truth of the various assertions in my paper to which Mr. Pedersen refers follows from the sampling interpretation.

I believe that Mr. Pedersen has misunderstood one aspect of the techniques described in my paper, probably because I adhered too closely to my constraint of a purely verbal description of the model and its related computational algorithms. The method for sampling paths of interest rates does not make the model stochastic at epoch 0 and deterministic at any or all future epochs. Nor does the sampling method render trading of bonds impossible in the model. At any time on any path, one has available by means of the simulation algorithm the full *yield curve*, not just the one-period rate of interest. Thus, one knows the prices of noncallable bonds of all maturities and one can trade a portfolio of bonds and establish a riskless hedge if desired.

Paths of interest rates from a given model can be simulated in many ways. A common approach is multinomial sampling in which N paths are generated from epoch 0 to epoch 1; for each of the N points at epoch 1, N paths are generated from epoch 1 to epoch 2; for each of the N^2 points at epoch 2, N paths are generated from epoch 2 to epoch 3; for each of the N^3 points at epoch 3, N paths are generated from epoch 3 to epoch 4; and so on. The structure of such a set of paths is a multinomial "tree" in which the number of paths grows exponentially as the number of time periods, making practical computations infeasible. It is better to sample paths as depicted in Figure 1 of Mr. Pedersen's discussion (with path crossings permitted, of course), but such sampling does *not* make the model deterministic beyond epoch 0. The model is stochastic beyond epoch 0 because a random number (or as many random numbers as there are stochastic factors) is needed at each epoch on each path

to extend the path to the next epoch. Because the full yield curve is simulated in an arbitrage-free manner at each epoch on each path, one does *not* need a "tree" emanating from each point to calculate the prices of all bonds at each point. This is the crucial concept underlying the practical implementation of the computational techniques that I have described.

Although I have not provided the mathematical details to make it easy, one can implement the algorithm described in Section 5.3 of my paper to construct a multifactor model of the term structure that satisfies the Heath-Jarrow-Morton (HJM) arbitrage-free conditions to which Mr. Pedersen refers in his discussion. [The article by Miller that is cited in my paper describes the approach. However, Miller's algorithm is based on spot rates, not forward rates, and develops negative forward rates with positive probability. Basing the processes explicitly on forward rates introduces some computational subtleties.] The resulting model can thus be characterized as an HJM model. Interestingly to Mr. Pedersen, I suspect, the model explicitly utilizes the two-step procedure described in his discussion. I intend to provide the technical details in a future communication, as I stated in Section 5.3.

Mr. Pedersen's question about what is an appropriate notion of an arbitrage-free model was also raised by Dr. Shiu during his reading of my paper prior to its being submitted for publication in the *TSA*. Obviously, I need to explain this important point more clearly and will try to do so here. The essential ideas are contained in Section 5.3 of my paper. To state the arbitrage-free condition as crisply as possible, a few definitions are helpful. An arbitrary epoch on an arbitrary interest rate path is chosen and referred to as the "initial point." Since full yield curves are simulated period by period under the model, there is no loss of generality in considering the yield curve obtaining at the initial point to be given— thus, it is referred to as "exogenously specified." The prices of all default-free zero-coupon bonds can be derived from the exogenously specified initial yield curve and are also considered to be "exogenously specified." Arbitrage-free adjustments of the drifts of the interest rate processes are defined by reference to these zero-coupon bonds.

- *Arbitrage-Free Adjustments.* The drift terms of the stochastic processes that govern the evolution of the yield curve from the initial point to the epoch one period ahead are adjusted at the initial point until the present value at the initial point of the expected one-period-ahead price of each zero-coupon bond is equal to its exogenously specified price at the initial point.

- *Arbitrage-Free Condition.* The stochastic interest rate model is said to be arbitrage-free when the arbitrage-free adjustments as defined above have been made at every epoch on every interest rate path.

Mr. Pedersen raises excellent points about determining the optimal exercise of options. First, I want to point out that the "evolution structure [of interest rates]" *is potentially available* at each epoch on each path—it is only *singly sampled* on each path after epoch 0 under my method for constructing paths. Second, I have written a paper that has been accepted for publication in Volume XLV of the *TSA* in which I demonstrate how to implement backward induction on a set of paths of the type depicted in Figure 1 of Mr. Pedersen's discussion. A specific numerical example involving an American put option is given in that paper. The valuation method is surprisingly accurate, even for relatively small sample sizes.

Finally, I encourage all readers of my paper to experiment with implementing the model and to plot the evolution of yield curves. I have implemented a four-factor HJM model that produces yield curves with the various shapes that are observed in the real world.

End Notes

1. A natural set of fundamental assets is all Arrow-Debreu securities. An Arrow-Debreu security pays $1 if a specified future state of the world occurs, and nothing otherwise. There are as many Arrow-Debreu securities as there are future states of the world. In a binomial model of interest rates, the two fundamental Arrow-Debreu securities defined at each two-pronged fork in the lattice are: an elementary put option that pays $1 if the up state occurs and $0 if the down state occurs, and an elementary call option that pays $1 if the down state occurs and $0 if the up state occurs. The role of the elementary options in valuing interest-rate-contingent cash flows is described in a paper by Jacob, Lord, and Tilley [15].

2. There are many ways that an expected present value could be computed for a stream of interest-rate-contingent cash flows. The intuitively appealing method described in the text above is the only one that is valid generally, but the proof is beyond

the scope of this paper. The essence of the proof for the case of binomial lattices (from which one can generalize to other situations) appears in Jacob, Lord, and Tilley [15] and utilizes the elementary put and call options described in footnote 1. Finally, I offer a helpful calculation reminder: On any path, be sure to use the one-period interest rate at epoch $t - 1$ to discount cash flow occuring at epoch t back to a present value at epoch $t - 1$.

3. The form of adjustment should be appropriate to the stochastic process that is modeled. A multiplicative adjustment of the type described in this section is valid for the lognormal process.

4. No distinction in notation is made in this paper between random variables and particular values they may assume, because the relevant interpretation is always clear.

5. Trying to cope with these difficulties, which arise more frequently when high yield volatilities, low mean-reversion strengths, or long-maturity bonds are used, quickly leads a model builder to the view that stochastic generators based on spot or forward rates would be highly preferred to those based on current-coupon yields.

6. For example, let Y be a random variable equal to the sum of n independent random variables X_i, $i=1, ..., n$, each of which has the same probability density function. Assume that the probability density function of X_t has zero mean and is symmetric. Then, regardless of sample size, the method of antithetic variates will calculate the median and all odd moments of Y exactly.

7. The method used by Balzer [1] on yield data for Australian government bonds will likely estimate the mean reversion level and strength parameters much more accurately than is possible by linear regression.

8. In using the equation given for the duration index *exactly* in the form written (that is, without some $(1 + i)$-type of factor multiplying the right-hand side of the equation), one should calculate the "shocked" price P' based on a parallel shift Δ in the *term structure*, not the yield curve, and one should express the spot rates defining the term structure as *forces of interest*. Otherwise, the duration index so calculated will not agree precisely with the conventional Macaulay-Redington definition.

9. Jetton, Merlin. "Interest Rate Scenarios," *TSA* XL, Part I (1989): 423-76.

10. Ibid., 460.

11. Mr. Pedersen, not a member of the Society, is a recipient of a Society of Actuaries Ph.D. Grant for his doctoral studies at Washington University, St. Louis.